

Joint and Actuator Design for Enhanced Stability in Robotic Force Control

by

Russell Duane Howard

B.S. California Institute of Technology (1977)

M.A. University of California, Berkeley (1979)

M.S. California Institute of Technology (1981)

SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September, 1990

© Massachusetts Institute of Technology, 1990.
All rights reserved.

Signature of Author _____
Department of Aeronautics and Astronautics, August 8, 1990

Certified by _____
Professor David L. Akin, Thesis Supervisor, Dept. of Aeronautics and Astronautics

Certified by _____
Professor Harold L. Alexander, Dept. of Aeronautics and Astronautics

Certified by _____
Professor Tomás Lozano-Pérez, Dept. of Electrical Engineering and Computer Science

Accepted by _____
Professor Harold Y. Wachman, Chairman, Departmental Graduate Committee

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

SEP 19 1990

LIBRARIES

Joint and Actuator Design for Enhanced Stability in Robotic Force Control

by

Russell Duane Howard

Submitted to the Department of Aeronautics and Astronautics on 8 August 1990
in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Abstract

Lightweight, high-transmission-ratio robot manipulators have not been able to achieve the theoretical force-control performance of a simple rigid-body manipulator. Avoiding instability on contact with the environment often requires severe limits on control gains and a soft covering for the end-effector (to limit the effective environmental stiffness). The results are low bandwidth, poor disturbance rejection, and degraded positioning accuracy.

For these manipulators, the worst problem is caused by dynamics *between* the actuator and the force sensor (the "non-colocated sensor" problem). These dynamics are due to drive compliance, link bending, actuator inductance, etc. With high-ratio transmissions, the dominant dynamic effect is usually drive compliance.

In this research, the force-control performance of a series of controller designs is studied using a compliant-drive manipulator model. A new design is suggested, based on the idea of decoupling the actuator from the joint, and controlling each with a separate feedback loop.

Analysis, simulation, and experiments are performed which show distinct advantages of the new approach over the best-performing reference design in the areas of stability, bandwidth, smoothness, and contact behavior. Although the new controller uses an internal model of drive compliance, its tolerance of modeling error is found to be excellent.

The principal cost of the new approach is an increase in actuator torque requirements, particularly when full advantage is taken of the improved bandwidth. Several means of reducing torque requirements are explored. A combination of actuator and joint torque command limits and filtering of the commanded reference trajectory succeeds in lowering peak torques by a factor of five, with no loss of performance.

Experiments performed on a single-joint manipulator testbed confirm the results of analysis and simulation, and validate the dynamic model used.

Thesis committee: Prof. David L. Akin, Chairman
Prof. Harold L. Alexander
Prof. Tomás Lozano-Pérez

Acknowledgements

All of the original research described in this document was performed in the MIT Space Systems Laboratory, funded by NASA Grant #NAGW-21.

The road to this degree has been long and challenging, with some interesting detours along the way. Many people have helped make the journey a successful and rewarding one.

Credit for any accomplishments must be shared with my family. This thesis is dedicated to my father, Lynn, who has taught me many things. Perhaps the most important is the value of persistence--a talent which I have found makes up for occasional lapses in other areas. My mother Doris has always encouraged me, and my sister Sharon's adventures continually broaden my horizons.

My feelings about MIT and NASA are mixed. I think M. A. Foster was correct in writing, "Individuals attain the power ... to create something; institutions, at vast and great labor, obtain finally the power to obstruct." Fortunately, I have encountered some outstanding individuals here and there within these institutions. Mel Montemerlo at NASA Headquarters and Dave Akin at MIT[†] are two such individuals whose vision and creativity have achieved a great deal. Their support has been essential to this project.

I am indebted to my thesis committee, Dave, Sandy, and Tomás, for their technical assistance and advice, and to Antonio Elias for his help in the early phases of my work. Craig Carignan labored with me to design and build much of the experimental hardware. Ping Lee, Al Shaw, Don Weiner, and Earle Wassmouth have always helped the Space Systems Lab make the most of its resources.

At a personal level, many have helped by sharing the ups and downs of MIT life. Janice Voss and her family have been particularly close friends. The last couple of years have been the most fun ever, partly due to some unique opportunities (courtesy of Bob Wolf, George Whittinghill, and the American Rocket Company), and partly due to an amazing collection of friends including Vicky Rowley, Rob Sanner, Ella Atkins, Matt Machlis, Killer Kowalski, Terry "Cheese Things" Fong, Jud "I dare you to be the last one" Hedgecock, and Nick the Cat.

I also gratefully acknowledge the support of my longtime companions Fluffy and Fang, and my really old friends Fred "ANIM8R" Kuentz and Pete Byrne (and the rest of the gang from S.I.).

[†] Now at the University of Maryland at College Park.

Table of Contents

Title Page.....	1
Abstract.....	3
Acknowledgements	5
Table of Contents.....	7
List of Figures	9
List of Tables	13
Chapter 1: Introduction	15
1.1 Motivation.....	15
1.2 Problem Statement.....	17
1.3 Thesis Overview.....	18
1.4 Contributions of this Thesis.....	20
Chapter 2: Force Control Performance.....	22
2.1 Transmission Ratio Effects.....	22
2.1.1 Effect of Transmission Ratio on System Mass	23
2.1.2 Side-Effects of Transmissions.....	27
2.2 Actuator Disturbance Forces	28
2.2.1 Friction	28
2.2.2 Cogging	29
2.2.3 Performance Effects.....	29
2.3 Backlash.....	31
2.4 Drive Compliance.....	33
2.4.1 Modeling Compliance for Control.....	34
2.4.2 Effects on Force Control Performance.....	37
2.5 Summary: Current Performance Barriers in Force Control.....	38
2.6 Manipulator Dynamic Model	40
2.6.1 Model Description.....	40
2.6.2 Controllability.....	41
2.6.3 Observability.....	43
Chapter 3: Force Controller Analysis	44
3.1 Force Control Objectives	45
3.1.1 Desired Closed-Loop Dynamics.....	45
3.1.2 Performance Criteria	46
3.2 A Brief Review of Control Theory	49
3.2.1 Stability Analysis.....	49
3.2.2 Second-Order Systems	51
3.3 The Forward Force Transfer Function.....	52
3.3.1 Definition.....	52
3.3.2 FFTF Evaluation for this Manipulator Model.....	54
3.3.3 Consequences for Stability.....	58
3.4 Force Controller Designs	60
3.4.1 The Return Force Transfer Function.....	61
3.4.2 Explicit Force Control	61

3.4.3	Stiffness Control	69
3.4.4	Stiffness Control with Inner Loop	77
3.4.5	Stiffness Control with Inner Loop and Filtering.....	85
3.5	Conclusions	92
Chapter 4:	Joint/Actuator Controller Design.....	94
4.1	Mechanical Decoupling of Joint and Actuator Dynamics	94
4.1.1	Smoothness Benefits.....	95
4.1.2	Contact Behavior Benefits.....	96
4.1.3	Other Performance Effects	96
4.1.4	Actuator Force Requirements.....	98
4.2	Joint/Actuator Control Loops	99
4.2.1	Controller Structure	99
4.2.2	Comparison with Force-Controlling Inner Loops.....	101
4.3	Performance Evaluation.....	107
4.3.1	Stability Range.....	108
4.3.2	Control Bandwidth.....	117
4.3.3	Smoothness	119
4.3.4	Contact Behavior	129
4.4	Conclusions: Performance Comparison Results	138
Chapter 5:	Simulation and Experiments.....	141
5.1	Simulation.....	141
5.1.1	Description of Simulation	142
5.1.2	Controller Comparison Example	143
5.1.3	A High-Speed Example.....	149
5.1.4	Effects of Modeling Error	152
5.1.5	Effects of Actuator Saturation.....	154
5.1.6	Effects of Filtering the Reference Trajectory.....	157
5.1.7	Low-Stiffness Example.....	159
5.1.8	Conclusions.....	165
5.2	Experiments	166
5.2.1	Description of Hardware	167
5.2.2	Results.....	171
5.2.3	Conclusions.....	184
Chapter 6:	Conclusions and Recommendations.....	185
6.1	Summary of Conclusions.....	185
6.2	Design Recommendations	187
6.3	Recommendations for Further Work.....	190
Appendix A:	Manipulator Mass Scaling Model.....	192
Appendix B:	Analysis of Backlash.....	198
B.1	Manipulator Model	198
B.2	Control Law	200
B.3	Stability Analysis	201
B.4	Derivation of Equivalent Linear System	203
B.5	Limit Cycle Prediction	206
B.6	System Simulation.....	208
Appendix C:	Controller Listings.....	212
References.....		217

List of Figures

2.1:	Manipulator Geometry	23
2.2:	Friction Models	28
2.3:	Backlash Input/Output Diagram	31
2.4:	Single-Joint Manipulator Model	40
3.1:	Feedback Block Diagram	49
3.2:	Block Diagram of Mechanism + Environment	52
3.3:	General Controller Block Diagram	53
3.4:	FFTF Magnitude Plot	56
3.5:	FFTF Phase Plot	56
3.6:	FFTF Magnitude; Stiff Environment	59
3.7:	Explicit Force Controller; Block Diagram	62
3.8:	Explicit Force Controller; RFTF Magnitude	63
3.9:	Explicit Force Controller; RFTF Phase	63
3.10:	Explicit Force Controller; Open-Loop Magnitude	64
3.11:	Explicit Force Controller; Open-Loop Phase	64
3.12:	Explicit Force Control; Closed-Loop Bandwidth	66
3.13:	Explicit Force Control; Force Disturbance Rejection	67
3.14:	Explicit Force Control; S/N Ratio	68
3.15:	Stiffness Controller; Block Diagram	70
3.16:	Stiffness Controller; RFTF Magnitude	71
3.17:	Stiffness Controller; RFTF Phase	71
3.18:	Stiffness Controller; Open-Loop Magnitude	72
3.19:	Stiffness Controller; Open-Loop Phase	72
3.20:	Stiffness Controller; Stability Tradeoff	73
3.21:	Stiffness Controller; Target Damping Tradeoff	74
3.22:	Stiffness Controller; Closed-Loop Admittance	75
3.23:	Stiffness Controller; Force Disturbance Rejection	76
3.24:	Stiffness Controller w/Inner Loop; Block Diagram	78
3.25:	Stiffness Controller w/Inner Loop; RFTF Magnitude	79
3.26:	Stiffness Controller w/Inner Loop; RFTF Phase	79
3.27:	Stiffness Controller w/Inner Loop; Open-Loop Magnitude	80
3.28:	Stiffness Controller w/Inner Loop; Open-Loop Phase	80
3.29:	Stiffness Controller w/Inner Loop; Stability Tradeoff	83
3.30:	Stiffness Controller w/Inner Loop; Closed-Loop Admittance	83
3.31:	Stiffness Controller w/Inner Loop; Force Disturbance Rejection	84
3.32:	Stiffness Controller w/Inner Loop + Filter; Block Diagram	86
3.33:	Stiffness Controller w/Inner Loop + Filter; RFTF Magnitude	88
3.34:	Stiffness Controller w/Inner Loop + Filter; RFTF Phase	88
3.35:	Stiffness Controller w/Inner Loop + Filter; Open-Loop Magnitude	89
3.36:	Stiffness Controller w/Inner Loop + Filter; Open-Loop Phase	89
3.37:	Stiffness Controller w/Inner Loop + Filter; Stability Tradeoff	90
3.38:	Stiffness Controller w/Inner Loop + Filter; Closed-Loop Admittance	90

3.39:	Stiffness Controller w/Inner Loop + Filter; Force Disturbance Rejection	91
4.1:	Joint/Actuator Controller Block Diagram.....	99
4.2:	Transmission Disturbance Model.....	103
4.3:	Transmission Models.....	105
4.4:	Decoupled-Drive FFTF Magnitude.....	109
4.5:	Decoupled-Drive FFTF Phase.....	109
4.6:	Joint/Actuator Controller; RFTF Magnitude	112
4.7:	Joint/Actuator Controller; RFTF Phase.....	112
4.8:	Joint/Actuator Controller; Open-Loop Magnitude	113
4.9:	Joint/Actuator Controller; Open-Loop Phase.....	113
4.10:	Joint/Actuator Controller; Stability Tradeoff.....	115
4.11:	Joint/Actuator Controller; Closed-Loop Admittance	118
4.12:	Force Disturbance Rejection Comparison	120
4.13:	Position Disturbance Rejection Comparison	122
4.14:	Reference Design; Joint-Position Response.....	125
4.15:	Reference Design; Actuator-Position Response	125
4.16:	Mechanically Decoupled Design; Joint-Position Response.....	126
4.17:	Mechanically Decoupled Design; Actuator-Position Response	126
4.18:	Joint/Actuator Design; Joint-Position Response.....	127
4.19:	Joint/Actuator Design; Actuator-Position Response.....	127
4.20:	Reference Design; Contact Force Response.....	133
4.21:	Joint/Actuator Design; Contact Force Response.....	134
4.22:	Joint/Actuator Design; Contact Force Response, Extended Interval	134
4.23:	Joint/Actuator Design w/Payload; Contact Force Response.....	136
5.1:	Reference Design; Simulated Position Trajectory	145
5.2:	Joint/Actuator Design; Simulated Position Trajectory.....	145
5.3:	Reference Design; Simulated Force Trajectory	146
5.4:	Joint/Actuator Design; Simulated Force Trajectory.....	146
5.5:	Reference Design; Simulated Control Input Trajectory	147
5.6:	Joint/Actuator Design; Simulated Control Input Trajectory	148
5.7:	Reference Design; Simulated Position Trajectory	150
5.8:	Joint/Actuator Design; Simulated Position Trajectory.....	150
5.9:	Reference Design; Simulated Force Trajectory	151
5.10:	Joint/Actuator Design; Simulated Force Trajectory.....	151
5.11:	Simulated Force Trajectory, +50 % \hat{k}_D , -50 % \hat{d}_D Modeling Error.....	152
5.12:	Simulated Force Trajectory, +50 % \hat{k}_D , +50 % \hat{d}_D Modeling Error.....	153
5.13:	Simulated Force Trajectory, -30 % \hat{k}_D , -50 % \hat{d}_D Modeling Error.....	154
5.14:	Simulated Position Trajectory; ± 2500 N Saturation Limit.....	155
5.15:	Simulated Position Trajectory; using Joint Force Saturation Limit.....	156
5.16:	Simulated Force Trajectory; using Joint Force Saturation Limit.....	157
5.17:	Simulated Force Trajectory; ± 1200 N Saturation Limit, Filtered X_{ref}	158
5.18:	Simulated Control Trajectory; ± 1200 N Saturation Limit, Filtered X_{ref}	158
5.19:	Reference Design; Simulated Position Trajectory, $\omega_0 = 2$ Hz	160
5.20:	Reference Design; Simulated Position Trajectory, $\omega_0 = 0.1$ Hz.....	160
5.21:	Reference Design; Simulated Force Trajectory, $\omega_0 = 0.1$ Hz.....	161
5.22:	Joint/Actuator Design; Simulated Position Trajectory, $\omega_j = 10$ Hz.....	162
5.23:	Joint/Actuator Design; Simulated Force Trajectory, $\omega_j = 10$ Hz.....	162
5.24:	Joint/Actuator Design; Simulated Force Trajectory, $\omega_j = 5$ Hz.....	163
5.25:	Joint/Actuator Design; Simulated Position Trajectory, $\omega_j = 1$ Hz.....	164

5.26:	Joint/Actuator Design; Simulated Force Trajectory, $\omega_j = 1$ Hz	164
5.27:	Experimental Apparatus, Top View	167
5.28:	Experimental Position Trajectory, Run 1	173
5.29:	Simulated Position Trajectory, Run 1	173
5.30:	Experimental Force Trajectory, Run 1	174
5.31:	Simulated Force Trajectory, Run 1	174
5.32:	Experimental Position Trajectory, Run 2	176
5.33:	Simulated Position Trajectory, Run 2	176
5.34:	Experimental Force Trajectory, Run 2	177
5.35:	Simulated Force Trajectory, Run 2	177
5.36:	Experimental Position Trajectory, Run 3	179
5.37:	Simulated Position Trajectory, Run 3	179
5.38:	Experimental Force Trajectory, Run 3	180
5.39:	Simulated Force Trajectory, Run 3	180
5.40:	Experimental Position Trajectory, Run 4	182
5.41:	Simulated Position Trajectory, Run 4	182
5.42:	Experimental Force Trajectory, Run 4	183
5.43:	Simulated Force Trajectory, Run 4	183
6.1:	Noise-Reduction Analogy	191
B.1:	2-Link Manipulator Model	198
B.2:	Limit Cycle Amplitude vs. Damping Ratio	206
B.3:	Limit Cycle Phase vs. Damping Ratio	206
B.4:	Limit Cycle Frequency Ratio vs. Damping Ratio	207
B.5:	Joint trajectory approaching limit cycle; $\zeta = 0.1$	208
B.6:	Joint trajectory in limit cycle; $\zeta = 0.1$	208
B.7:	Joint angles vs. time; $\zeta = 0.1$	209
B.8:	Joint trajectory in limit cycle; $\zeta = 2.0$	209
B.9:	Joint angles vs. time; $\zeta = 2.0$	210

List of Tables

2.1:	RMS Example Parameters.....	25
2.2:	Dexterous Example Parameters.....	26
3.1:	Control Analysis Parameters.....	55
3.2:	Sample Stiffness Table	69
4.1:	Controller Comparison Parameters.....	107
4.2:	Stick-Slip Simulation Parameters.....	124
4.3:	Performance Parameter Comparison.....	138
5.1:	Controller Comparison Simulation Parameters	143
5.2:	Low-Stiffness Simulation Parameters	159
5.3:	Experimental Apparatus Parameters.....	170
5.4:	Experimental Parameters	171
A.1:	Inertial Scaling Parameters	195
B.1:	Limit Cycle Parameters	210

Chapter 1: Introduction

1.1 Motivation

There is a growing need for robotic manipulators to perform tasks which would be too costly or dangerous for humans. Examples occur undersea, in space, and on the factory floor. Robots are potentially faster, stronger, and more precise than humans, and less subject to fatigue. For space applications (of particular interest here) they are also likely to require less support equipment and maintenance.

Why then are they not commonplace in space and installed in every factory? What has prevented them from fulfilling their potential?

Part of the answer lies in the type of manipulation current robots are capable of. The first generation of commercial robot manipulators found several niches in earthbound industry, notably for welding and painting tasks in mass production. These successful applications were characterized by two features: predictability and repetition. The objects handled had known shape and mass properties. Robot workspaces included fixtures to make the geometry consistent and predictable.

Predictability is necessary because these early robots were position-controlled machines with little or no external sensing. They were good at following programmed paths and rejecting any disturbances which would cause them to deviate. This is desirable behavior for many industrial tasks. However, a much broader class of useful tasks require the path to vary in response to interaction forces which are not known in advance. This is known as *compliant motion*^[26,36], and the *compliance* of a manipulator is a specification of the amount and direction of position deflection produced by an imposed force.

A standard example is the “peg-in-hole” task. To insert a rigid peg in a close-fitting hole using position control requires extremely accurate knowledge of the hole’s location and precise control of the peg. Small errors cause the peg to jam, preventing completion of the task and generating large, potentially damaging contact forces. In practice, the necessary precision can be achieved with a specialized machine, but rarely with a general-purpose robot. Alternatively, jamming can be avoided by even a low-precision manipulator

if the robot reacts properly to interaction forces as a human does in performing this type of task.

Compliant motion is a better strategy than pure position control for tasks requiring motion while in contact with the environment. This is particularly true for doing work on the environment (in the physics sense, of applying a force during motion). In the factory, compliant motion is essential for such economically useful tasks as assembly of components.

In space as well, the capability of compliant motion would make a robot much more effective by increasing the range of tasks it could perform. Robotic manipulators in space will be used for exploration, assembly, maintenance, and repair. Since transportation costs to orbit are extreme (around \$4800/kg to low Earth orbit^[18]), and few tasks are likely to be standardized and repetitive, versatility is a key feature for space manipulators.

Compliant motion has been the focus of a great deal of research^[5,26,36]. Two types of manipulators have been developed to meet this need: those with “passive” compliance, and those with “active” force control.

In the first type, compliance comes from the mechanical design of the manipulator or (usually) the end-effector, which deflects in a predetermined way when forces are encountered. Standard position control drives the joints to follow a nominal trajectory, while the mechanical compliance moderates interaction forces and aligns the payload to facilitate the desired task.

Quite a bit of work has been done to define the types of passive compliance which are useful, and ways to achieve these in a mechanism. Whitney’s Remote-Center Compliance (RCC)^[53] is one such concept that has already found commercial acceptance. These devices (and other advances such as machine vision) are making the current generation of industrial robots more capable and less dependent on task predictability. However, the passive compliance of a manipulator must be set up for a particular task and is not easily changed. This limits the versatility of an individual robot and is acceptable in many factory situations, but a drawback for space applications.

The alternative to passive compliance is active force control, which promises to improve versatility by allowing the compliance to change between (and during) tasks. With active force control, compliance is achieved by modifying joint torques in response to the sensed interaction forces. A variety of ways of implementing this have been considered^[2,13,21,27,28,34,40,42,47,54,57]. In this thesis, a *force-controlled manipulator* is a robot with force sensing at or near the endpoint and a closed-loop force-feedback control scheme. The controller is used to impose the commanded compliance on the manipulator.

Motion in response to sensed forces follows a user-specified law; for example, the target compliance could require the manipulator to react like a damped spring when pushed.

The initial attempts to implement active force control met severe performance barriers^[39,54]. Recent studies^[16,26,51] have identified some of the principal causes of instability and error. For the lightweight, high-transmission-ratio manipulators discussed here the worst problem is caused by dynamics *between* the actuator and the force sensor (the “non-colocated sensor” problem). These dynamics include transmission compliance and link bending.

Some of these problems are eased considerably with the direct-drive actuators adopted by several researchers^[2,57]. Unfortunately, the adverse effect on total system mass of using direct drive makes it impractical for most space applications.

1.2 Problem Statement

The specific problem addressed by this thesis is: **Develop and evaluate new design principles to improve the force-control performance of manipulators with transmission dynamics.** Transmission dynamics, broadly defined, includes all physical effects which alter the rigid-body relationship of actuator motion to joint motion. Here, the dominant effect is assumed to be compliance in the drive train.

A secondary goal is to shed further light on the underlying mechanics of manipulators and force control algorithms.

There are two reasons to focus on manipulators with transmission dynamics. First, when direct-drive manipulators can be used, their performance is already quite good. Second, due to the intended space applications, low overall mass is essential. As described in Section 2.1, this requirement leads to designs with transmissions which multiply torque and allow remote actuator installation. Designers use gearboxes, ball screws, harmonic drives, chains, cables, torque tubes, etc., which introduce significant dynamics of their own. This necessitates the specific treatment of drive compliance and actuator-axis disturbance forces such as friction.

Transmissions, of course, are useful in many earthbound applications as well. Lower mass and remote actuators often mean less bulk, allowing slimmer manipulators with faster response and greater dexterity. The results of this research should be of general interest to the robot-design community.

1.3 Thesis Overview

Much progress has been made recently in understanding the observed limitations on force control performance^[16,26,51]. The problem now becomes to translate this insight into design strategies which improve performance.

The earliest suggestions along these lines followed directly from the performance analyses. For example, if limit cycles were found to result from stiction (a type of friction), stiction should be reduced; if instability resulted from drive compliance, the drive should be stiffened, etc.

Although such recommendations are sound, there are limits to this approach. Stiction and drive compliance, for instance, are costly to reduce and already minimal in a good mechanism. Moreover, reducing them may be incompatible with design goals such as low overall size or mass.

More fundamental progress will come from control and design strategies tailored to cope with the known disturbances. Instead of simply minimizing an acknowledged source of error, an understanding of it may allow its effects to be ameliorated by a change of linkage geometry or the method of control. Familiar examples of this are the use of dither to reduce friction effects without reducing the friction itself, or the addition of an inner control loop to better reject actuator disturbance forces. Since the order of complexity of the mechanism is often increased, such changes must be made with an understanding of the overall system dynamics.

The objective of this thesis is to build upon recent analytical work characterizing the limitations on force control performance, by recognizing these sources of error and instability and taking them into account in a design strategy. This is an open-ended problem since as performance levels increase more dynamic effects will become important.

The key ideas developed in this thesis are the use of Forward and Return Force Transfer Functions for force controller analysis and tuning, the mechanical decoupling of small-scale actuator dynamics from joint dynamics, and the separate control treatment of actuator and joint.

Quite a bit of background information is helpful in understanding the problems posed and solutions suggested here. The second and third chapters of this thesis are an attempt to provide this by laying out the problem in detail, and defining the current state of the art. No new strategies are introduced here, but existing ones are illuminated in new ways.

Chapter 2, “Force Control Performance,” defines the problem explicitly. The relevant physical phenomena are discussed and their effects on performance described. Substantial review of the current literature is included in this section, although the analysis of backlash and the manipulator mass scaling models are original to this thesis. A simple dynamic model of a single manipulator joint with drive compliance is adopted for subsequent analysis and simulations.

In Chapter 3, “Force Controller Analysis,” a set of performance criteria is defined to evaluate improvements. These are stability, bandwidth, smoothness, and contact behavior. A method for analyzing force control algorithms is introduced, using Forward and Return Force Transfer Functions. This method is applied to a series of controller designs drawn from the literature to explore how they work, what their limitations are, and to indicate how to overcome these.

A force-controlling manipulator can be thought of as a loop: control input is processed through manipulator and workpiece dynamics to give output force, and this force is sensed and processed through controller dynamics to give further control input. Previous studies in the literature have concentrated on the first half of the loop, adding and subtracting various modeled physical phenomena to determine their effects on performance. The analysis in this chapter focuses instead on the return half of the loop. Assuming a fixed physical model, various controller features are added to determine their effects on performance. A standard example quantifies the performance improvements.

Chapter 4, “Joint/Actuator Controller Design,” introduces a new manipulator design approach based on the separate treatment of joint and actuator motion; hence the term “Joint/Actuator Controller.” The two inertias are mechanically decoupled by *adding* a linear elastic element to the transmission, and the controller design forms separate (nested) feedback loops around each. The addition of a joint position sensor is also required.

The new controller includes a model of drive compliance, and is intended to overcome the bandwidth limit set by the drive resonance frequency. Stability range should also improve, allowing lower target stiffnesses to be achieved. Mechanical drive decoupling is intended to improve disturbance rejection and ensure the validity of the controller’s internal model.

Section 4.3 compares the Joint/Actuator design to the best-performing design of Chapter 3. The comparison is made using the performance criteria defined in Section 3.1.2, and the results show improvement in every category.

Chapter 5, “Simulation and Experiments,” extends this analysis to the time domain through numerical simulations which include nonlinear effects such as finite sampling rate, sensor and control resolution, actuator saturation, dry friction, contact discontinuity, etc.

Also described are the results of experiments comparing the reference design to the new ideas on a single-joint laboratory testbed.

Simulation is used to confirm the behavior of the controllers predicted in Chapters 3 and 4 to discover the effects of modeling error and to investigate techniques for improving the practicality of the Joint/Actuator design by reducing actuator force requirements. Experiments are performed to validate the manipulator dynamic model used throughout this thesis and to directly demonstrate the stability advantages of the new controller.

Chapter 6 reviews the major conclusions. Also presented are recommendations for implementing the ideas discussed, and topics for further work.

1.4 Contributions of this Thesis

This section summarizes the most important contributions of this research:

- A detailed investigation is performed of the problem of force control with a compliant-drive manipulator. An analytical framework using Forward and Return Force Transfer Functions is developed and used to study and compare existing force-control algorithms.
- The analysis suggests a new approach: the Joint/Actuator controller design with mechanical drive decoupling. The new design is shown to permit operation above the actuator-drive resonance frequency ω_D and offers lower attainable target stiffnesses than conventional designs, with improved disturbance rejection.
- A stability criterion is derived for the Joint/Actuator controller, enabling the designer to select the control gains necessary to achieve the desired target stiffness and bandwidth with a given environment. It shows that the stability limit is now due to the limited bandwidth ω_a of the inner loop, rather than the drive resonance at frequency ω_D . This has two effects which improve stability over conventional designs: the inner-loop bandwidth is generally higher than ω_D and the inner-loop damping ratio is much greater than that of the drive resonance.
- A study of manipulator contact behavior characterizes the initial (open-loop) response during a collision with the environment, indicates the closed-loop bandwidth necessary to improve this, and examines the roles of drive compliance and the inertia *outboard* of the force sensor (including payload). It is found that drive compliance typically has little effect, but that payload mass places a fundamental limit on the ability of a controller to limit impact forces and bouncing.

- A study of ways to reduce actuator force requirements with the Joint/Actuator controller reveals the utility of filtering unneeded high frequencies out of the commanded position trajectory. It is also shown that limits (in software) on commanded joint force assist the recovery from momentary saturation-induced instability, providing a good safety feature for this type of controller.
- An experimental testbed was constructed. Experiments confirm the expected stability improvements of the new design, and validate the dynamic model and simulation procedure.
- A number of design recommendations are made. Some of these are specific to the Joint/Actuator design, such as the desirability of locating the friction-introducing part of the transmission inboard of the decoupling springs (e.g. installing the gearbox on the actuator rather than the joint). Others are more generally applicable, such as the benefit of minimizing structural mass outboard of the force sensor.
- Aspects of the Joint/Actuator design can be generalized to broader classes of mechanisms. The modeling approach and controller structure allow stable operation above the natural frequency of transmission dynamics, without requiring a model of the environment; other effects could similarly be modeled and compensated for. Also, the principle used to improve disturbance rejection can be abstracted as follows: Use a software pre-filter (in this case the drive model) to boost input signals, and a matching inverse mechanical post-filter (the decoupled drive dynamics) to restore the desired output signal (in this case interaction force) while attenuating disturbances. This resembles such electronic noise-reduction schemes as the *Dolby™* system, and is discussed more fully in Section 6.3.

Chapter 2: Force Control Performance

This chapter characterizes some of the physical phenomena relevant to force control performance in order to identify the most immediate challenges.

The effects of transmission ratio are described first. The ratio, an important parameter of any transmission, is defined as the number of turns of the input shaft required to produce one turn of the output shaft. The claim was made in the introduction that a high-ratio transmission is required to meet the objective of low system mass. In the first section of this chapter, this claim is justified with a simple mass-scaling model. When a high-ratio transmission is used the effects of actuator-axis disturbance forces (those acting on the motor shaft), backlash, and drive compliance become important, and these are discussed in subsequent sections.

In the final two sections of the chapter, the dominant physical effects are summarized and a dynamic model is defined for the class of mechanisms of interest. This model is used throughout the rest of the thesis as a basis for analysis and simulation, and its accuracy is evaluated experimentally in Chapter 5.

2.1 Transmission Ratio Effects

For the design engineer, the ability to specify the transmission ratio between each actuator and the manipulator joint it drives is an important one. It allows the designer to adapt actuator characteristics to the expected loads.

Direct-current electric motors provide a simple, approximately linear means of converting electrical signals to force and motion, and are commonly used as actuators in robots. Most are capable of high speeds (thousands of rpm), and their bulk, weight, and cost rise rapidly with torque capability. Robotic applications, however, typically require low speeds for a fraction of a revolution, with high torque. Transmissions allow high-speed low-torque actuators to meet the low-speed high-torque demands of manipulators.

2.1.1 Effect of Transmission Ratio on System Mass

When the designer is free to select a transmission ratio, higher performance can be achieved for a given system mass. The effect is quantified in this section.

A mass-scaling model is developed in Appendix A to compare two types of serial-link, revolute-joint electric manipulators. One design uses direct-drive actuators and the other uses transmissions which provide torque multiplication and allow the actuators to be installed remotely from the joint. Displacement of the actuator is used to reduce the moment of inertia of each link.

Manipulator Scaling Model

The model allows the calculation of total manipulator mass given the desired payload mass, tip acceleration, link length, and a parameter related to link stiffness. The geometry shown in Figure 2.1 is assumed.

In the arm with transmissions the actuators have been displaced from the joints they drive but are still installed on the same link. This allows some reduction of inertia with minimum added mechanical complexity. Some manipulators remove all of the actuators to the base, giving the maximum reduction of inertia.

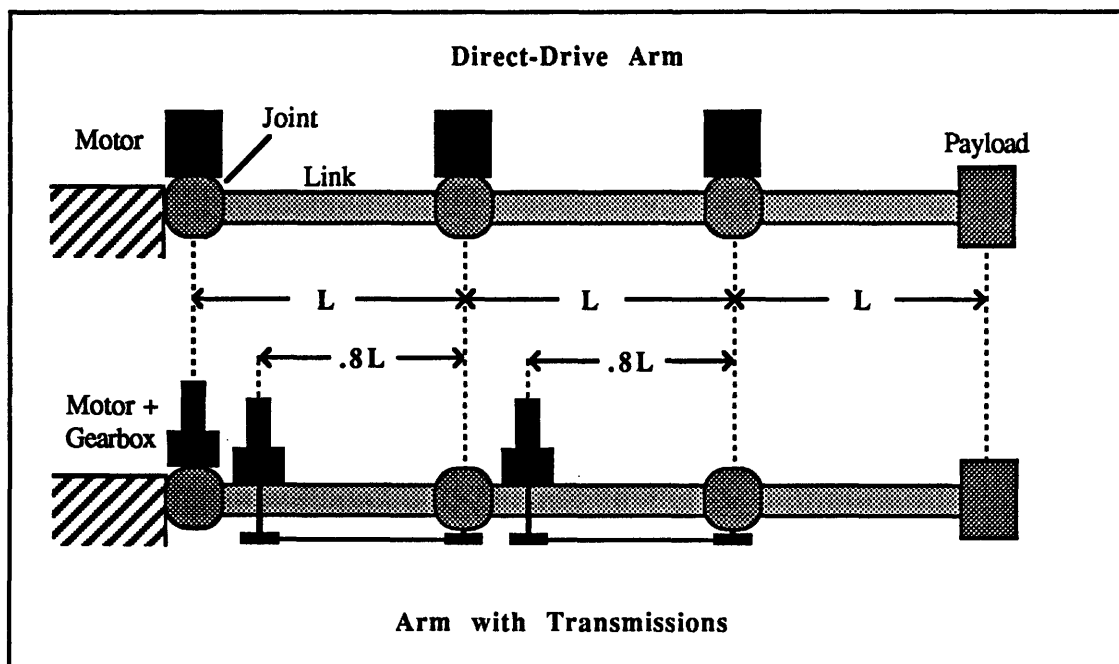


Figure 2.1: Manipulator Geometry

The independent (input) parameters are:

$L \equiv$ link length

$A \equiv$ required tip acceleration capability

$M_p \equiv$ payload mass

$M_{L3} \equiv$ mass of third link

The dependent (calculated) parameters are:

$M_{L1}, M_{L2} \equiv$ mass of links 1, 2

$M_i \equiv$ mass of motor i

$M_{Gi} \equiv$ mass of gearbox i

$M_T \equiv$ total manipulator mass

$\dot{\omega}_i \equiv$ angular acceleration of link i

$J_i \equiv$ moment of inertia of arm about joint i

$J_{Mi} \equiv$ rotor inertia of motor i

$J_{Ti} \equiv$ total inertia seen by motor i

$T_i \equiv$ torque applied at joint i

$T_{Mi} \equiv$ torque applied by motor i

$N_i \equiv$ transmission ratio at joint i

The equations for calculating these parameters are developed in Appendix A.

Manipulator Examples

Two examples are calculated to compare the overall mass of direct-drive vs. geared manipulators.

RMS-Equivalent Manipulator

An interesting comparison is for performance equivalent to the Space Shuttle RMS (Remote Manipulator System) which is 18 meters long, weighs 413 kg, has a payload mass of 30000 kg, tip acceleration of .0025 m/sec², and gear ratios of 2000:1.

The actual RMS doesn't quite fit the manipulator design template of Figure 2.1. It has two major links rather than three and, although it uses geared transmissions, the actuators are located directly at each joint. What will be done here is first to design a geared 3-link manipulator which fits the template and matches the RMS values for total mass, total length, payload mass, and tip acceleration, and which uses the actual RMS gear ratios. In this procedure total mass is an independent parameter and M_{L3} becomes a derived parameter. This is done by choosing a trial value for M_{L3} , calculating M_T and iterating

until the total mass is that of the real RMS. Essentially, this is using the real RMS mass to set the link stiffness parameter in our model.

This gives a RMS-equivalent geared manipulator which fits the template. The equations developed above are then used to design a direct-drive version with the same performance.

The assumptions and results for this RMS-equivalent example are shown below:

<u>Specifications</u>		
<u>Parameter</u>	<u>Value</u>	
L	6 m	
M_p	30000 kg	
A	.0025 m/sec ²	
M_{L3}	27 kg (derived from matching M_T for geared version to real RMS value)	
<u>Results</u>		
<u>Parameter</u>	<u>Geared Value</u>	<u>Direct-Drive Value</u>
M_1	15.3 kg	478 kg
M_2	11.9 kg	360 kg
M_3	7.7 kg	222 kg
M_{L1}	243 kg	245 kg
M_{L2}	108 kg	108 kg
M_T	413 kg (specified)	1440 kg (derived)
Mass Ratio: 3.48		

Table 2.1: RMS Example Parameters

Since the geared version was intended to match the real RMS, equation (A.13) was not used in this case. Instead, the actual ratios $N_i = 2000$ were substituted. This leads to some error, as the algorithm outlined above still assumes that only half of the motor torque goes toward accelerating the arm (the rest accelerates the rotor inertia). With these ratios the effective rotor inertia is less than in the impedance-matched case, so more torque would be available for accelerating the link. The effect on system mass is small and does not change the nature of the results; in fact, a more accurate calculation (requiring iteration) would show a larger advantage for geared systems.

The model predicts that for a 3-link version of the RMS, direct drive would increase mass by over 1000 kg for equivalent performance, reducing Shuttle payload capacity by the same amount, at an approximate cost per mission of \$5,000,000.

Dexterous Manipulator

The RMS is basically a crane for gross positioning of large objects. At the other end of the manipulation spectrum would be a smaller unit for tasks requiring dexterity: .5 meter links, payload mass 10 kg, tip acceleration 1 m/sec². The assumptions and results for this dexterous manipulator example are shown below:

<u>Specifications</u>		
<u>Parameter</u>	<u>Value</u>	
L	.5 m	
M _p	10 kg	
A	1 m/sec ²	
M _{L3}	1 kg	
<u>Results</u>		
<u>Parameter</u>	<u>Geared Value</u>	<u>Direct-Drive Value</u>
M ₁	.875 kg	33.1 kg
M ₂	.643 kg	19.5 kg
M ₃	.414 kg	10.0 kg
M _{L1}	10.7 kg	16.8 kg
M _{L2}	4.2 kg	5.2 kg
M _T	17.85 kg	85.6 kg
Mass Ratio: 4.80		

Table 2.2: Dexterous Example Parameters

The model predicts a geared version would have a mass of 18 kg and the direct-drive version 86 kg, giving a mass ratio of 4.8.

The overall effect on system mass will be even more than these figures indicate. The more massive direct-drive arm will require more power at higher current levels to follow the same trajectory, so additional mass is needed for power supply components (e.g. batteries). If the arm is mounted on a free-flying platform, more fuel and structure

will be required. All of this extra mass must be lifted to orbit. Thus, the decision to use direct-drive actuators on a space manipulator would be an expensive one.

2.1.2 Side-Effects of Transmissions

A real transmission has several effects on manipulator joint dynamics besides the desired torque multiplication^[1,48].

The effect of motor armature inertia is multiplied by N^2 . This has the benefit of making the overall manipulator dynamics less configuration-dependent. The inertia matrix is closer to diagonal and constant, so the nonlinear rigid-body dynamics of a multi-DOF manipulator lose importance relative to the dynamics of each joint.

Drive compliance is increased, as if a lightly-damped spring had been placed between the actuator and the link. This is caused by gear tooth bending, shaft bending, cable stretching, etc. Friction forces within the actuator are multiplied along with the actuator force: Coulomb friction and stiction by the transmission ratio N , and viscous friction by N^2 . More friction may arise within the transmission itself, from gear tooth sliding, intermediate shaft bearings or tendon sheaths^[24], for example. If gears are used backlash and cogging are present^[1]. These effects can seriously degrade the performance of a manipulator using transmissions. The following sections describe this in more detail.

2.2 Actuator Disturbance Forces

A single manipulator link can be modeled as having an inertia associated with the actuator (the actuator node) and an inertia associated with the joint (the joint node). This is the basis for the dynamic model adopted later for comparative analyses. If transmission dynamics are present between them, these nodes may follow quite different trajectories. Disturbance forces which act directly on the actuator node are particularly important in manipulators with transmissions because their effect on joint motion is amplified by the transmission ratio. Unlike forces acting at the joint node, they can be corrected immediately if sensed or predicted accurately. This section describes the most commonly-encountered actuator disturbance forces and their effects on performance.

2.2.1 Friction

“Friction” refers to any forces arising from the contact of two surfaces which oppose their relative motion. This includes viscous forces from lubricant wetting the surfaces, as well as dry friction.

Friction forces are primarily a function of the relative velocity of the surfaces and the normal force, and are observed to fall into three categories^[11,50]: viscous friction, Coulomb friction, and stiction. In viscous friction the force is proportional to the velocity; in Coulomb friction the force depends only on the sign of the velocity, and stiction forces act only at zero (or very small) velocity.

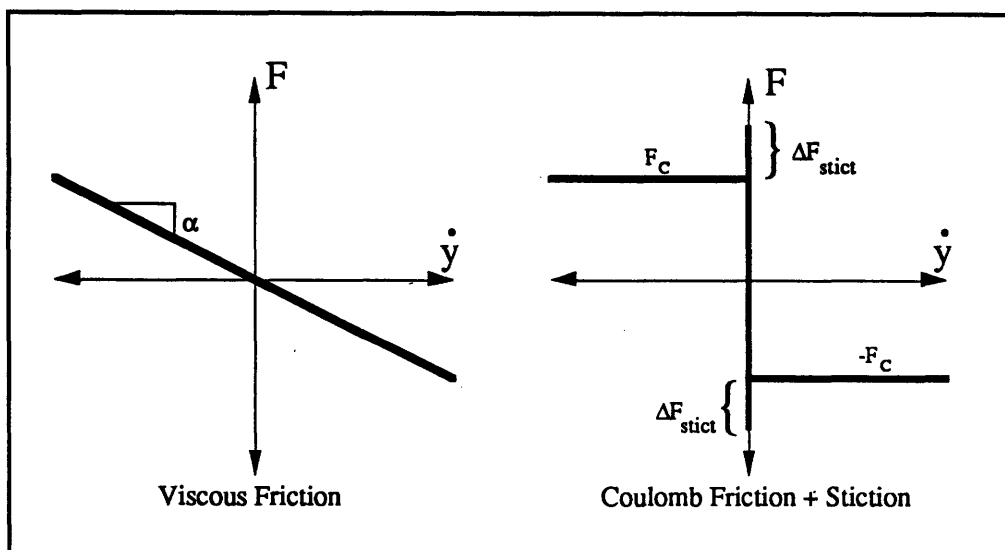


Figure 2.2: Friction Models

These are idealized models of actual friction forces. Canudas et al.^[11] performed a series of experiments on a prototype manipulator which found forces conforming to each of these types, but they also found significant time and temperature variation in the parameters, and the slope of the viscous friction component depended on the direction of motion.

2.2.2 Cogging

Disturbance forces which depend on actuator position rather than velocity are described here as cogging forces. This includes the uneven torque provided by a gear train due to the small changes in geometry as the teeth mesh, and ripple torques in the motor due to changing magnetic field geometry as the rotor turns. These disturbances are related to the number of teeth on the gears or number of poles in the motor. Eccentricity and offsets in shafts and bearings also contribute to cogging^[1].

With rotary actuators, cogging forces are periodic. Since they act on the actuator node and intermediate gear meshes, the disturbance frequency is proportional to transmission ratio for a given joint velocity.

Measurements made by the author on experimental manipulators in the laboratory indicate that cogging forces may amount to 5 % or more of the maximum torque output of a typical industrial-quality DC motor/gearbox actuator.

2.2.3 Performance Effects

Of the disturbance forces catalogued above, only viscous friction admits a linear model. Viscous friction appears as the familiar damping or dissipation term in linear differential equations. The usual effect of actuator damping is to increase stability margins at the expense of bandwidth.

The methods of nonlinear analysis^[20,37] seldom allow generalizable conclusions about the effects on stability or performance of a given disturbance. In many cases, one must fall back on simulations of the particular situation of interest. Nonetheless, recent work by Townsend and Salisbury^[50] sheds light on the qualitative effects of dry friction (Coulomb friction and stiction) on force control. Their analysis pertains to a simplified manipulator model, with an actuator capable of instantaneously attaining a commanded velocity. In practice this would correspond to a manipulator with a low-mass actuator, a low transmission ratio, and a high-gain inner control loop based on actuator velocity. A compliant transmission connects the actuator to the inertial joint node, upon which the friction acts. Note that this is a different problem than discussed above; the disturbance

forces act on the joint node, not the actuator node. This distinction diminishes, however, in the near-inertialess case which is the basis for several of their results.

They conclude that Coulomb friction introduces a time delay in response, but can increase system stability for low amplitudes of force error, leaving high-amplitude stability limits unchanged. In the near-inertialess case, stiction always resulted in limit cycles--a periodic stick-slip motion of the joint. The amplitude of the limit cycle was proportional to ΔF , the stiction force.

Other researchers^[40] have noted limit cycles in actual manipulator behavior which they attribute to an interaction between stiction and an integrator in the control loop.

2.3 Backlash

Backlash is a commonly encountered nonlinearity which can lead to limit-cycle instabilities in an otherwise stable system. It arises in standard geared transmissions from the necessary gap between teeth of mating gears^[1]. When drive direction is reversed, free motion occurs between the input and output shafts as this gap is traversed. Figure 2.3 illustrates the relationship between input angle and output angle.




Figure 2.3: Backlash Input/Output Diagram

“Anti-backlash” gears avoid this by splitting one gear (e.g. the input gear) into two parallel gears connected by a torsion spring. One of these gears contacts one side of the teeth of the output gear and the other contacts the other side, so that no gap exists. This eliminates backlash, but at the same time increases friction and drive compliance. Since transmissions for robotic force control are often designed for minimum friction and compliance, backlash can still be a problem.

A nonlinear control analysis was undertaken to investigate the effects of backlash on a serial-link rotary-joint manipulator. The analysis (detailed in Appendix B) shows:

- Backlash inserted between the actuator and control sensor on an otherwise stable manipulator system introduces a single stable limit cycle into its closed-loop free-space behavior.
- The limit cycle frequency is proportional to but less than controller bandwidth ω_c . Increasing the controller damping ratio ζ reduces the limit cycle frequency.
- Limit cycle amplitude is strongly dependent on controller damping. As $\zeta \rightarrow 0$, amplitude $\rightarrow \infty$; as $\zeta \rightarrow \infty$, peak-to-peak amplitude decreases asymptotically to a minimum at the backlash magnitude b .
- When limit cycle amplitude significantly exceeds backlash magnitude ($\zeta < 1$) motion is smooth and periodic; at smaller amplitudes the higher harmonics become important and motion becomes quasi-periodic.

These results apply whether the backlash physically occurs in the actuator or in the control sensor. Note that if there are gravity forces or if the manipulator is in contact with its environment, enough bias force may be present to suppress actuator backlash, although sensor backlash would remain. In more complicated systems with multiple control loops these results still apply, with the relevant ω_c and ζ being those associated with the innermost loop containing the backlash.

The third point above is important to note for the controller analysis in subsequent chapters. It implies that the amplitude of a backlash-induced limit cycle cannot be reduced below a minimum value by increasing the gains. The forces involved may increase, but the position amplitude remains nearly constant. This is an essential nonlinear aspect of backlash.

In linear control analysis, disturbances are usually formulated as *forces*, and disturbance rejection calculated as the transfer function from this force to output error. Backlash limit cycles, however, more closely resemble a constant-amplitude *position* disturbance of the actuator (assuming reasonable damping of the actuator control loop). It is therefore more relevant to examine the transfer function from actuator *position* to output error.

2.4 Drive Compliance

Transmissions inevitably add some elasticity, or compliance, to the drive train of a manipulator. Indeed, elastic elements are essential to the function of some types such as harmonic drives.

Drive compliance can degrade performance in two ways. In position control applications, the sensors used for control can measure either actuator shaft angles or joint angles. If actuator angles are used for control, any compliance between the actuator and joint introduces a load-dependent error in actual payload position. This sets an upper limit on the effective stiffness of the manipulator, no matter how high the control gains are set. Conversely, if joint angles are used for control, it becomes a classic “non-colocated” sensor problem. The compliance introduces destabilizing phase shift at high frequencies between sensor and actuator, limiting the allowable gains and, again, the achievable stiffness.

In force-controlled manipulators, compliant dynamics between the force-sensing wrist and the actuators cause instability at high gain^[15,54]. The worst case occurs in contact with a stiff environment, when force feedback effectively becomes very high-gain non-colocated position feedback^[2,25,40,41]. This problem can be more troublesome than in position control because the effective gain depends on the environment. Oddly enough, this sets a *lower* limit on achievable stiffness, because in this case the force feedback gains are higher for a low target stiffness (since the motion required in response to a sensed force is greater).

Control theory has given the designer tools which work well for controlling rigid manipulators, as experiments with direct-drive attest^[2,57]. Once it is understood that drive compliance can cause stability problems in manipulators, it is logical to formulate a plant model which includes the drive compliance. The principles of control theory might then be expected to yield robust, high-performance control laws as in the rigid case. Unfortunately, as the next section describes, this has not proven to be easy.

2.4.1 Modeling Compliance for Control

In order to use the powerful methods of linear control theory, the plant equations, written in input/output form, must be linear and uncoupled.

Manipulator dynamics contain many nonlinear terms. However, the dynamics of a rigid manipulator can be linearized and decoupled by nonlinear *static-state* feedback. In other words, an equation of the following form can be written:

$$\underline{u} = \underline{c}(\underline{x}) + D(\underline{x}) \underline{a} \quad (2.4.1)$$

where:

- \underline{u} = control variables (vector); in this case actuator torques
- \underline{x} = state variables (vector); in this case joint angles and rates
- \underline{a} = reference variables (vector); in this case joint accelerations
- \underline{c} = nonlinear function of states (vector); in this case Coriolis and centripetal forces
- D = nonlinear function of states (matrix); in this case the inertia matrix

The important feature of equation (2.4.1) is the linear, uncoupled relationship between \underline{a} and \underline{u} , which allows the application of linear control methods. Controller design using this equation is known as the resolved-acceleration^[33] or computed-torque^[19] method.

A manipulator with drive compliance cannot be linearized in this way^[35], due to additional internal state variables. Each inertial node contributes two state variables--its position and velocity. The link inertias now move independently of the actuator inertias, so each link brings with it four state variables instead of two.

What can be used instead is nonlinear *dynamic-state* feedback^[12]. This allows equations of the following form to be written:

$$\dot{\underline{z}} = \underline{b}(\underline{x}, \underline{z}) + B(\underline{x}, \underline{z}) \underline{a} \quad (2.4.2)$$

$$\underline{u} = \underline{c}(\underline{x}, \underline{z}) + D(\underline{x}, \underline{z}) \underline{a}$$

where: \mathbf{z} = new state variables (vector); actuator angles and rates
 \mathbf{x} = the original state variables (vector); joint angles and rates
 \mathbf{a} = reference variables (vector); in this case not the joint accelerations
 \mathbf{b}, \mathbf{c} = nonlinear functions of states (vector)
 \mathbf{B}, \mathbf{D} = nonlinear functions of states (matrix)

The first equation contains the dynamics of the drive compliance. The most important practical difference between the dynamic-state formulation and the static-state one is the fact that \mathbf{a} no longer represents the joint accelerations. This is a disadvantage, since the overall goal is to control joint behavior.

The difference can be understood by examining the physical system. With a rigid drive, it is clear that at any point in time the acceleration of a link will have a linear relationship with the actuator torque, because this torque is applied directly to the link inertia. The slope and offset of this relationship vary with the system state, as given by equation (2.4.1), but the relationship is always linear. With drive compliance, however, a torque applied to the actuator inertia will have a linear relationship with the *actuator* acceleration, not the link acceleration.

If, for example, drive compliance is represented by a simple spring, the displacement of the spring gives the force transmitted from actuator to joint. The *displacement* of the actuator then has a linear relationship with the *acceleration* of the joint. Since the actuator torque is linearly related to the actuator acceleration, it is also linearly related to the fourth derivative of the joint angle. The reference vector \mathbf{a} in this case would be the fourth derivative of the joint motion, rather than simply joint acceleration.

The example of a three-link manipulator with a model of drive compliance was investigated in [12]. In this case, the linearizing, decoupling transformation (2.4.2) used makes the dynamical system equivalent to three chains of six integrators each, instead of the double integrators of the rigid-body case. In other words, the reference \mathbf{a} turned out to be a vector of sixth derivatives of joint angles, rather than accelerations. It is conjectured in [12] that similar results apply to any arm with their model of drive compliance.

Unfortunately, while they are linearized and decoupled, it is hard to design a real controller with these equations. To produce the usually-desired second-order behavior by controlling sixth derivatives of joint position would require measurement or estimation of the fourth and fifth derivatives. Noise and unmodeled dynamics make this impractical.

Approximations to the exact dynamics (2.4.2) are often considered. One approach is to linearize the original nonlinear equations about an operating point, i.e. a given state \mathbf{x}_0 . Experiments have shown that good position-control performance can then be

obtained with a linear controller in the neighborhood of \underline{x}_0 [23]. For serial multiple-link manipulators, however, the dynamics vary greatly over the workspace. This restricts the range of validity of the linearized model.

Another way to approximate the exact dynamics uses singular perturbation theory [29]. The addition of drive compliance is considered as a (small) perturbation to the rigid-body manipulator dynamics. The equations can be separated into “slow” dynamics and “fast” dynamics. A slow controller is designed for the unperturbed (rigid) system using a model of the form (2.4.1), and a “corrective control” is added to make the slow dynamics behave like the rigid-body system. This corrective control is not explicitly solvable, but can be approximated to any desired order [46]. A fast controller is used to stabilize the fast dynamics, ensuring that trajectories converge to the slow dynamics [35,44]. In this case, the fast controller damps out drive oscillations. While this has the potential of making a compliant-drive manipulator behave like a rigid one, the approximation that allows separation into slow and fast dynamics limits the overall closed-loop bandwidth (that of the slow controller) to below the natural frequency of the compliance.

Another problem with many higher-order model approaches is that they require an equally accurate model of the environment, since in force control the environment is “in the loop” and the coupling is strong through the force sensor. This would add much complication to a working manipulator system, and robustness might be poor.

In summary, modeling compliance is a good idea but has been difficult to take advantage of in practice. Success that has been achieved in position control of flexible-link manipulators has required a good model and sophisticated control, and even then is conditionally stable [9,10].

The exact problem is very difficult; what is needed is a simple approximation to it which allows a practical controller design. The approach taken in this thesis makes such an approximation, and enforces it by redesigning the mechanism to ensure that the approximation is sufficiently accurate.

2.4.2 Effects on Force Control Performance

Transmission compliance contributes extra phase shift to the open-loop transfer function at high frequencies^[16,51]. This leads to instability as the closed-loop bandwidth of the system (including the environment it is coupled to) approaches the natural frequency of the compliance.

The documented effects of such a stability limit are to limit bandwidth and enforce a tradeoff between stiffness of the environment and the achievable closed-loop compliance of the manipulator^[16,54]. A rigid environment requires stiffer manipulator behavior. The recommendations presented so far in the literature are twofold:

- Through mechanical design, raise the natural frequency of the drive dynamics as high as possible, and
- Add compliance between the force sensor and the environment, to limit the effective environmental stiffness seen by the manipulator.

To obtain a large improvement by following the first recommendation requires a direct-drive manipulator. As shown in Section 2.1 this is not always practical. The second recommendation is more easily followed--by installing compliant coverings on grippers, for example. There are limits here as well; adding too much compliance makes for a bulky end-effector and degrades positioning precision.

2.5 Summary: Current Performance Barriers in Force Control

Constraints on robot mass and bulk make the use of high transmission ratios essential for many applications.

Lightweight, high-ratio manipulators have not been able to achieve the theoretical force-control performance of a simple rigid manipulator, which is approached by some direct-drive machines. This failure is due to higher-order dynamic effects which lead to instability and disturbances which lead to error.

The dynamic effects include force sensor flexibility, motor inductance, link bending, drive compliance, and delay due to digital sampling. All of these introduce phase shift between the force sensor signal and the actuator control signal at high frequencies, causing instability when the closed-loop bandwidth is too high.

Disturbances include sensor noise, friction, stiction, backlash, and cogging at the actuator axis, at the joint axis, and within the transmission.

The most important dynamic effects are those occurring at the lowest frequencies; these are the immediate barriers to improving closed-loop bandwidth and stability. The assumption of a high transmission ratio helps determine which effects are dominant. Measurements of high-ratio manipulators in use in this laboratory give the following typical values for the frequencies at which different effects become important (defined here as contributing at least 45° of phase shift):

Motor Inductance	125 Hz
500 Hz Sampling Rate	62.5 Hz
Force Sensor Flexibility	30 Hz
Link Bending	25 Hz
Drive Compliance	3 Hz

Of these, the most immediate challenge is clearly drive compliance.

For disturbances, the most important are simply those with the greatest effect on endpoint errors (force and position). With typical transmission ratios > 100, the effect of actuator-axis force disturbances far outweighs that of joint-axis force disturbances. Friction and cogging will be considered actuator-axis force disturbances and the limit-cycling due to backlash will be considered an actuator-axis position disturbance.

The effect of sensor noise will not be explicitly considered here, although it may be significant in some cases. Digital sensors (encoders or resolvers) are commonly used for position measurements and are essentially noise-free, although they have finite resolution.

Force sensors, on the other hand, are analog devices and susceptible to electrical and thermal noise. Excessive force sensor noise could limit force-control performance. However, the position response of a force-controlled manipulator to a sensed force is usually intended to be inertial; i.e. the response to a force is acceleration. Noise in the force sensor is thus integrated twice before it appears as position error. This diminishes the effect of high-frequency, zero-bias noise. Low-frequency noise, or calibration drift, is not attenuated as much and should be limited by careful design.

2.6 Manipulator Dynamic Model

For the analysis in subsequent chapters, a dynamic model is required to represent force-controlled manipulators. The model should be as simple as possible while retaining the dynamics being studied. Much research on robot modeling has been done^[1,2,5,16,41,48,51]. Eppinger's work^[16] in particular has shown the utility of single-joint lumped-parameter models in understanding the performance of real manipulators in force control.

For multi-DOF force control, each axis can be considered to be autonomous^[47]. This just reflects the fact that a Cartesian-coordinate end-effector stiffness can be transformed into joint-coordinate stiffnesses using the inverse Jacobian (if one avoids singularities). Accordingly, most of the remainder of this thesis assumes a single-joint manipulator model. The important dynamic effects described above can be included, and all of the performance criteria proposed can be quantified and compared using such a model.

2.6.1 Model Description

The following diagram illustrates the single-joint manipulator model:

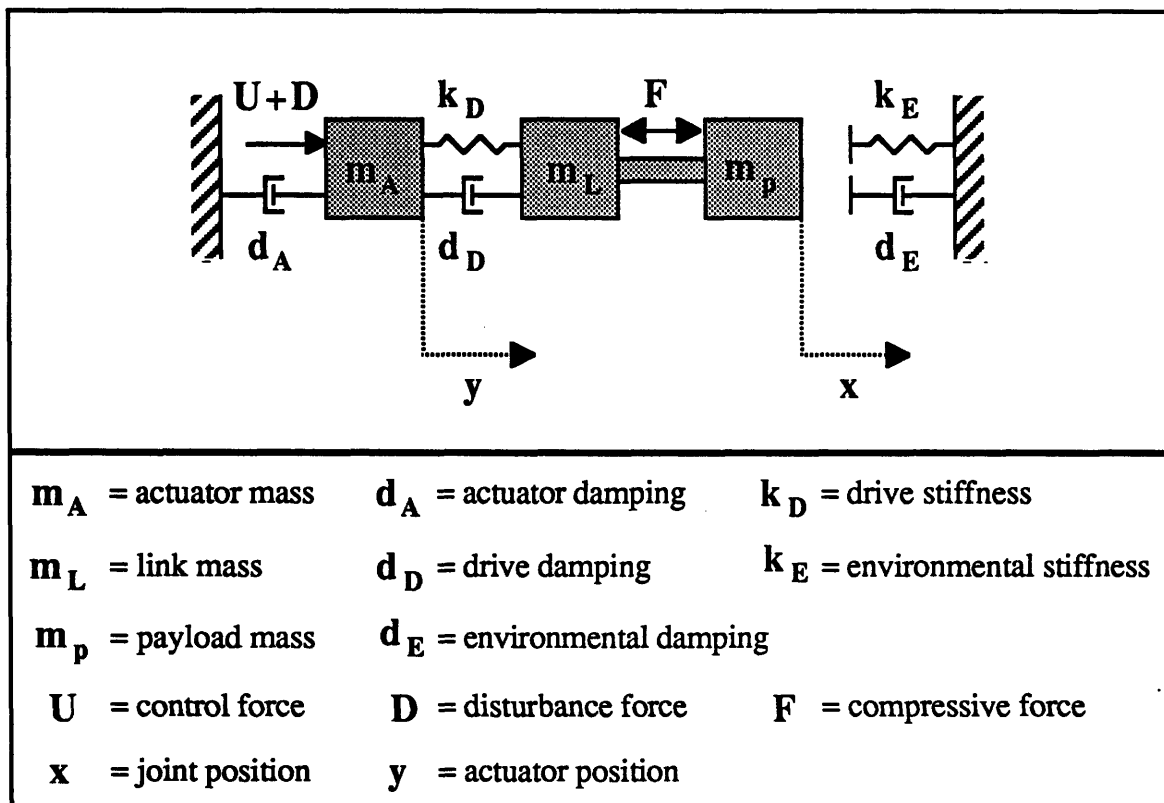


Figure 2.4: Single-Joint Manipulator Model

The disturbance force D includes Coulomb friction and stiction. The actuator is assumed to accept a force command U . Compressive interaction force F is measured by a (rigid) sensor at the “wrist,” with m_P representing the mass (including payload) outboard of the sensor. The link itself is assumed rigid. Position coordinates are the joint position x and the actuator position y . The environment is represented by a spring and damper for $x > 0$, and no effect for $x \leq 0$.

The assumption of a rigid force sensor is an approximation which may fail in practice. An important effect of flexibility in the force sensor would be to limit the effective environmental compliance to the value of the compliance in the sensor. Experiments in Chapter 5 show that when this effect is taken into account, the model above provides a good representation of actual hardware.

The dynamic equations of this system are:

$$m_A \ddot{y} = U + D - d_A \dot{y} - d_D (\dot{y} - \dot{x}) - k_D (y - x)$$

$$m_L \ddot{x} = d_D (\dot{y} - \dot{x}) + k_D (y - x) - F \quad (2.6.1)$$

$$F = m_P \ddot{x} + [d_E \dot{x} + k_E x]_{x>0}$$

In this model, the usually rotary motions of joints and actuators have been abstracted to linear motions. The physics is equivalent, but terminology changes. Actuator and joint *torques* become *forces*, and actuator and joint *angles* will be referred to as *positions*.

2.6.2 Controllability

Controllability is a property of a linear system which implies that it is possible to drive the system from any initial state to any final state in finite time, with appropriate control inputs. A system without this property may have oscillatory modes which cannot be damped out by control signals, or other undesirable behavior.

If the manipulator model adopted above is uncontrollable, efforts to improve its performance are doomed from the start. It would be necessary to redesign the modeled system to ensure control effectiveness over all the states.

Controllability is easily checked. The system can be written in state-space form $\dot{\underline{x}} = \underline{A}\underline{x} + \underline{B}u$, where \underline{x} is the vector of states and u is the control input. \underline{A} , \underline{B} and \underline{x} are shown in equation (2.6.2).

$$\underline{A} \equiv \begin{bmatrix} -\frac{d_D + d_E}{m_L + m_P} & -\frac{k_D + k_E}{m_L + m_P} & \frac{d_D}{m_L + m_P} & \frac{k_D}{m_L + m_P} \\ 1 & 0 & 0 & 0 \\ \frac{d_D}{m_A} & \frac{k_D}{m_A} & -\frac{d_A + d_D}{m_A} & -\frac{k_D}{m_A} \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (2.6.2)$$

$$\underline{B} \equiv \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \quad \underline{x} \equiv \begin{bmatrix} \dot{x} \\ x \\ \dot{y} \\ y \end{bmatrix}$$

The controllability matrix is then defined as $[\underline{B} | \underline{A}\underline{B} | \underline{A}^2\underline{B} | \dots | \underline{A}^{n-1}\underline{B}]$. Control theory asserts that the system is controllable if the rank of this matrix is equal to n , the dimension of the state vector^[43].

In this case the controllability matrix is a 4 X 4 matrix whose determinant must be nonzero. After some algebra performed with the aid of the *MACSYMA*TM computer program, this determinant is found to be:

$$\det(\text{controllability matrix}) = \frac{(m_L + m_P)k_D^2 + d_D^2k_E - d_Dd_Ek_D}{m_L^3 + 3m_Pm_L^2 + 3m_P^2m_L + m_P^3} \quad (2.6.3)$$

Since m_P and m_L are always positive, only the numerator is of concern. Note that when $d_E = 0$ as in free space or against a purely spring-like surface, the numerator is positive and the system is controllable. The same is true if drive damping $d_D = 0$. Since both of these parameters are small compared to the other terms in practice, the system is theoretically controllable in all situations of interest.

This assures us only that control signals can exist to make the manipulator behave as desired. No guidance is given as to what sensors or feedback law may be used to generate them.

2.6.3 Observability

A related property to controllability is observability. A system is *observable* if every state can be exactly determined from measurements of the output over a finite interval of time.

If the manipulator model is not observable with the sensors available, the controller cannot drive it to the desired state. Also, if transfer functions are to be used in controller analysis, observability is a requirement for accurate representation of the system.

The equation $\underline{y} = \underline{C}\underline{x}$ represents the sensor readings as a function of state. The observability matrix is then defined as $[\underline{C}^T | \underline{A}^T \underline{C}^T | \underline{A}^{T2} \underline{C}^T | \dots | \underline{A}^{Tn-1} \underline{C}^T]$, and must be of rank n for observability^[43].

For the manipulator model used here, the states are $x, \dot{x}, y,$ and \dot{y} , as shown in (2.6.2). Clearly if all of the states are directly measured (\underline{C} is the identity matrix), the system is observable. In many manipulators only the actuator states y and \dot{y} are directly measured, along with the interaction force F . Some algebra with MACSYMA indicates that this is still observable. In fact, if the environment dynamics are included and payload mass is nonzero, the system is in principle observable using only the interaction force sensor F .

Chapter 3: Force Controller Analysis

The first section of this chapter specifies what is meant by force control in this thesis, and proposes four criteria for good performance. This is followed by a discussion of analysis techniques for force controllers, including development of the concept of Forward and Return Force Transfer Functions. The remainder of the chapter is devoted to exploration of the properties of a series of controllers drawn from the literature, to reveal the source of their limitations and discover ways to overcome these.

Analysis of the controllers here is similar in methodology to work done by other researchers^[16,51], in the use of simple models and linear control techniques. However, those efforts focused on exploring the behavior of a single assumed (very simple) controller with models of a variety of physical phenomena, to determine the effects of these phenomena on force control. The approach here is the converse; to select a single representative physical model (aided by their findings), and explore the effects of a variety of controller designs.

This work also serves to verify that the simple dynamic model adopted at the end of Chapter 2 can account for the performance limitations seen in various experiments. The most important of these is instability upon contact with a rigid environment, which has required either severe lowering of control gains^[8] or the installation of a very compliant covering for the end-effector^[7]. Another observed feature of past implementations of force control is poor disturbance rejection^[34,54], leading to uneven behavior and large errors.

The controller analysis in this chapter confirms the results of other research^[16,51] that these effects can be traced to the destabilizing nature of feedback from a non-colocated force sensor, and goes on to show the utility of position-based inner loops and filtering of the force feedback. These features have been demonstrated before^[22,34], but the discussion here clarifies why they work, and examples allow quantitative performance comparisons of different controllers for the same system. The best of the controllers is selected as a performance baseline for evaluating the improvements described in the next chapter.

3.1 Force Control Objectives

Chapter 2 described the interfering effects of various physical phenomena on the behavior of force-controlled manipulators. This section looks at the problem from the user's point of view to determine what performance features are most important. Taking into account these goals and the nature of the disturbances, performance criteria are picked to guide improvements and enable comparisons with other work.

3.1.1 Desired Closed-Loop Dynamics

To date, there are no commercially-available manipulators using active endpoint force control, yet a great deal of research has been done on this topic in the last decade. High hopes are held out for a wide range of assembly tasks^[4,21,25,26,28,36,40,55] currently beyond the capabilities of position-controlled manipulators.

“Force control” itself means many things to many people. Once a manipulator can react to measured forces, the question becomes how *should* it react. The range of behaviors that ultimately may be desired of manipulators is unknown, but the alternatives proposed so far in the literature are fairly simple and some have been shown to facilitate experimental tasks. The hybrid control of Raibert and Craig^[40], Salisbury's stiffness control^[42], damping (also called accommodation) control^[13,52], Hogan's impedance control^[21,27], and various types of explicit force control are among them. Many of these are reviewed by Whitney^[54]. In general, the desired behavior in these schemes is to exert a commanded force (force error is corrected immediately), or simulate a specified damper (force error produces velocity response) or damped spring (force error produces position displacement).

In explicit force control, only the force trajectory is commanded. In stiffness (or damping) control a reference position trajectory is given, along with the commanded spring (damping) constant. “Impedance” control requires the manipulator endpoint to follow second-order dynamics about the reference trajectory, with a commanded mass as well as stiffness and damping. These distinctions blur a bit in practice; all controllers have limited bandwidth and ultimately behave like inertial objects, so can be considered special cases of impedance control whether the target mass is implicit or explicit.

For the purposes of this thesis, the pragmatic view is taken that the user will want as much bandwidth as he/she can get. In other words, the user would not command a higher target mass than required for stability and acceptable tracking. The bandwidth then will be implicit, and the user will command only stiffness and damping about a reference

position trajectory. It is still sometimes useful to break this down into two classes of controller: those which are primarily springs (with some damping), and those which are primarily dampers (possibly with some “spring” to give self-centering and prevent divergence).

3.1.2 Performance Criteria

The following are the author’s attempt to make precise and measurable those qualities of performance which are i) of value to the user, and ii) difficult to achieve with current manipulators. Many researchers have formulated similar lists which overlap this one to some degree^[48,51].

Stability Range

Stability, in some sense, is required for every operating condition encountered as the manipulator performs its tasks. The strictest control theory definition--uniformly, asymptotically, stable in the large (UASIL) is not necessary; bounded instabilities such as limit cycles are acceptable if of small enough amplitude.

Divergent instability is often observed to limit the range of a performance-related figure of merit, such as k_E/k_T , the ratio of environmental stiffness to commanded manipulator stiffness, in the case of a spring-like controller contacting a spring-like environment^[54]. The particular figure of merit depends on the task at hand and the structure of the controller used. For stiffness controllers, the range of stiffness k_T and damping d_T achievable with a given environment is a good measure, and range of stability offers a way to compare the performance of different controller implementations.

Control Bandwidth

Control bandwidth is defined here as the maximum frequency at which the manipulator system can track the desired dynamics. Note that the desired dynamics may include a very soft spring and thus exhibit low-bandwidth response to external forces, and yet the control bandwidth could remain high. This is simply a case of a high-bandwidth controller tracking a trajectory which lacks high-frequency content. The full bandwidth may not be needed for command following, but is available for disturbance rejection.

In many situations bandwidth is directly related to productivity, hence economic viability. Regardless of the desired dynamics, higher control bandwidth improves performance. Bandwidth must be evaluated in conjunction with stability range, since there is usually a tradeoff between the two.

Smoothness

The motion of early force-controlled manipulators has been described as “hesitant,” and “jerking”^[34]. This can be traced to the influence of disturbances like friction and backlash. Both give rise to limit cycles, and friction causes stick-slip behavior during motion.

Since small position errors can generate large contact force errors, force tracking may be poor. Even when maximum tracking errors are small, these disturbances can add unwanted higher harmonics to the trajectory. This in turn can excite undesirable dynamics in the payload or environment. In a task such as grinding, trajectory smoothness is reflected in the quality of the surface produced. For the assembly of delicate parts, smoothness may mean the difference between success and failure.

A logical way to evaluate smoothness is to examine the transfer function from disturbance forces at the actuator to contact forces at the manipulator endpoint. Desirable features would be low gain amplitude within the controller’s bandwidth and rapid gain rolloff with frequency. Simulation must be used to check performance with specific nonlinear types of disturbance, e.g. stiction.

Disturbances such as actuator limit cycles are better seen as position disturbances than as force disturbances. It is important in these cases to look at the transfer function from actuator *position* to endpoint force.

A closely related parameter is the available force dynamic range, defined as the ratio between the maximum usable force a manipulator can apply at its tip and the minimum increment of force it can reliably apply. This is a useful dimensionless measure of the versatility of a manipulator, indicating the range of tasks it can perform^[51]. A high ratio of these forces implies a good combination of strength and delicacy.

The maximum usable force is determined by actuator capacity multiplied by transmission ratio and link geometry factors, minus the peak value of disturbance forces which must be countered. The minimum force unit can depend on a variety of factors: force sensor resolution, actuator force resolution, limit cycle magnitude, etc. In a manipulator with force feedback, it can be evaluated using the disturbance transfer functions and estimates of the disturbance magnitudes.

Contact Behavior

A recognized difficulty in manipulation is the “contact” problem^[28,39]. When a manipulator contacts the environment, its dynamics change from an open kinematic chain to a closed one. Ideally, one moment there is relative velocity and the next there is none. The design problem is to ensure that the kinetic energy is dissipated without damage and to provide a rapid, smooth transition to constrained motion. Since with a rigid payload and environment the initial collision may take only microseconds, the amount and distribution of passive (mechanical) compliance and inertia in the manipulator play an important role before the controller can react.

Relevant measures of contact performance are peak impact force and the time from initial contact to continuous contact, both of which should be minimized.

3.2 A Brief Review of Control Theory

Since stability is critical to performance, some concepts from control theory will be used frequently in the analysis and are reviewed here. Currently available microprocessor speeds are high and likely to increase further, so the analysis is performed in continuous time rather than discrete time. This implies that sampling rates can be set fast enough not to be the factor limiting performance. Continuous-time analysis is clearer and more straightforward since the equations representing physical systems are derived as differential, rather than difference, equations.

3.2.1 Stability Analysis

In classical linear control theory, dynamics are often described in terms of *transfer functions*. Following [38], the transfer function of a linear time-invariant system is defined to be the ratio of the Laplace transform of the output to the Laplace transform of the input, under the assumption that all initial conditions are zero. A transfer function is an accurate representation of a system if and only if the system is controllable and observable^[43]. A *sinusoidal transfer function* is a transfer function with s , the Laplace variable, replaced by $i\omega$, where ω is a (real) frequency in radians/second, and $i \equiv \sqrt{-1}$.

In a system with feedback, stability depends on the locations of the poles of the closed-loop transfer function. For stability, they must all have negative real parts. There are many ways of checking this, given a system model: Nyquist diagrams, Routh test for polynomials, iterative root-solvers, etc. For several reasons, the method used here is to look at the Bode plot of the open-loop transfer function.

Assume a system is described by the following block diagram:




Figure 3.1: Feedback Block Diagram

The closed-loop transfer function is:

$$\frac{X(s)}{X_{\text{ref}}(s)} = \frac{G(s)}{1 + G(s)H(s)} \quad (3.2.1)$$

The Nyquist stability criterion states^[38]: If the open-loop transfer function $G(s)H(s)$ has k poles in the right-half s plane, then for stability the $G(s)H(s)$ locus, as s traces out the modified Nyquist path in the clockwise direction, must encircle the $-1 + 0i$ point k times in the counterclockwise direction. The modified Nyquist path is simply a contour in the s -plane which encloses the right half-plane without passing through any poles or zeroes of the function $G(s)H(s)$. For systems with minimum phase open-loop transfer functions (no poles or zeroes in right half plane) $k = 0$; no encirclements are allowed.

A Bode diagram consists of two graphs: one is a plot of the logarithm of the magnitude (in db) of a sinusoidal transfer function; the other is a plot of the phase angle; both are plotted against the frequency ω in logarithmic scale. The Bode diagram of a system's open-loop transfer function shows the behavior on the portion of the Nyquist path that traces the imaginary axis. For a minimum-phase system, the stability criterion requires that the phase at the gain crossover frequency (the frequency at which gain = unity) be greater than -180° .

Phase margin is defined as the amount of additional phase lag at the gain crossover frequency required to bring the system to the verge of instability. *Gain margin* is the reciprocal of the magnitude of $G(s)H(s)$ at the frequency where the phase angle is -180° . Both can be easily read from the Bode diagram.

Bode diagrams of the open-loop transfer function offer several advantages over other methods of stability assessment. They give an indication of *relative* stability through gain and phase margins. The frequency-response approach itself often provides physical insight. When transfer functions are multiplied, their Bode plots simply add; this allows the effect on stability of particular physical phenomena to be identified. Since modeling accuracy is usually good at low frequencies and poor at high frequencies, the sensitivity of a system's stability to modeling error can be estimated. There is also the practical advantage that Bode plots can be obtained experimentally for an unmodeled system.

3.2.2 Second-Order Systems

A common factor in transfer functions is second-order in s : $(As^2 + Bs + C)$. A few comments reviewing the behavior of such a factor are in order. In physical terms, the coefficient A is analogous to mass, B to damping, and C to spring stiffness. Substituting $i\omega$ for s , it can be seen that the real part goes to zero at the “natural frequency”:

$$\omega_n \equiv \sqrt{\frac{C}{A}} \quad (3.2.2)$$

in radians/second. The magnitude of the function at this resonant frequency is determined solely by the imaginary part, and is $B\omega_n = B\sqrt{\frac{C}{A}}$. Note that this magnitude is proportional to the damping coefficient B . The dimensionless “damping ratio” is defined as:

$$\zeta \equiv \frac{B}{2\sqrt{AC}} \quad (3.2.3)$$

A system is called “underdamped” if $0 < \zeta < 1$, “overdamped” if $\zeta > 1$, and “critically damped” if $\zeta = 1$.

When an underdamped second-order factor appears in the denominator of a transfer function, the result is a resonance peak in magnitude. The peak occurs at the “damped natural frequency” $\omega_{dn} \equiv \omega_n \sqrt{1 - \zeta^2}$, and its magnitude is equal to $\frac{1}{\zeta \sqrt{4 - 3\zeta^2}}$. An

underdamped system exhibits oscillatory behavior in the time domain.

In this thesis, simple PD control loops are often discussed. These controllers have a proportional gain C and a derivative gain B . When used to control the position of an inertia A , the closed-loop dynamics are second-order, as above. If the inertia is known, specifying the natural frequency and damping ratio is equivalent to specifying the gains, and vice-versa. Equations (3.2.2) and (3.2.3) can be used to compute these quantities from the gains and the following gives the inverse relationship:

$$\begin{aligned} C &= A \omega_n^2 \\ B &= 2 \zeta A \omega_n \end{aligned} \quad (3.2.4)$$

Since natural frequencies and damping ratios are often more intuitively meaningful than gains, this method of description will be used frequently. For inner position loops the implied inertia A will be the actuator inertia m_A , and for outer position loops the implied inertia will be the link inertia m_L . The natural frequency of a control loop will also be referred to as its bandwidth.

3.3 The Forward Force Transfer Function

Since the purpose of this chapter is to investigate controller designs, a particular way of dividing up the system is chosen which isolates controller dynamics in the feedback transfer function $H(s)$, and makes the forward transfer function $G(s)$ dependent only on the mechanism and environment dynamics, which remain fixed. The forward function $G(s)$ reveals features of the plant which must be dealt with by any controller, and the properties of the controllers themselves can be compared through their return (feedback) functions $H(s)$.

3.3.1 Definition

The system model described in Section 2.6 can be depicted by the following block diagram:

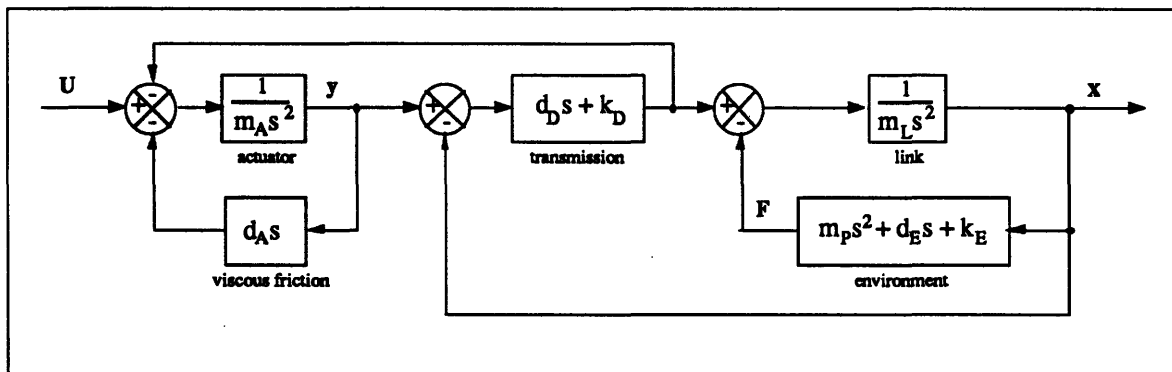


Figure 3.2: Block Diagram of Mechanism + Environment

A controller closes the loop. A general form for controllers is assumed to be:

$$u = A(s) F + B(s) x + C(s) y + R(s) r \quad (3.3.1)$$

where r is a command input--a reference force or position.

The controller block diagram is then:

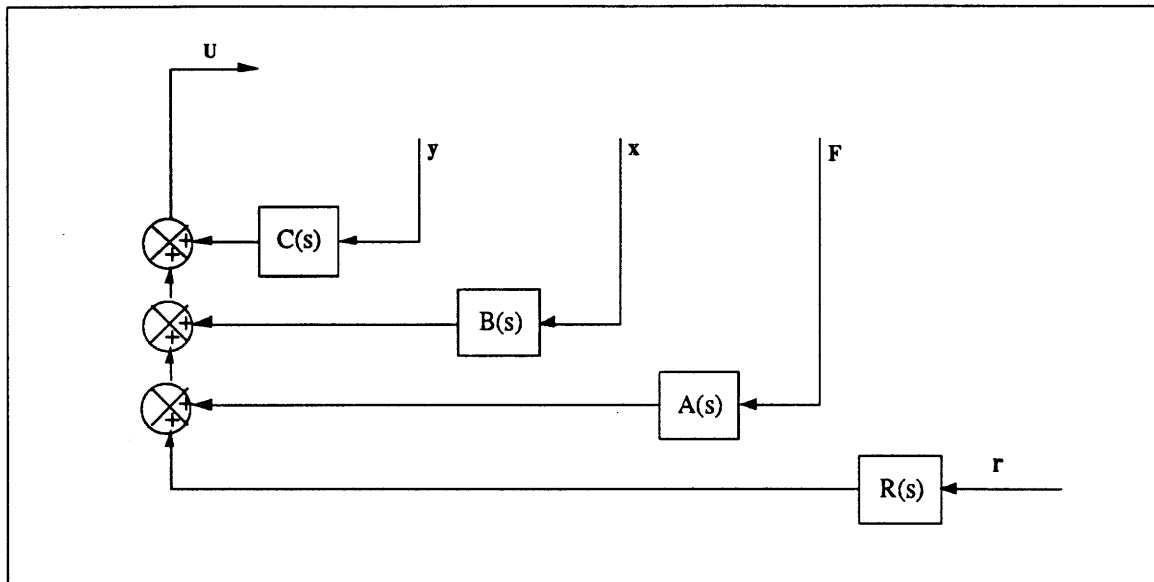


Figure 3.3: General Controller Block Diagram

Through block diagram reduction techniques^[38], these transfer functions can be combined into a system with a block diagram of the form of Figure 3.1. Bode plot stability analysis can then be carried out on the reduced system.

For the forward transfer function, consider the loop to be broken between the controller and the actuator, and the output to be the force F ; choose $G(s) \equiv \frac{F(s)}{u(s)}$. This is dimensionless and easy to interpret physically, as the transmission function from actuator force to measured force at the wrist, when the manipulator is in contact with the environment. From this alone it can be guessed that $G(s) \rightarrow 1$ at low frequencies, and any drive compliance will cause the magnitude to roll off and the phase to lag at high frequencies. Since $G(s)$ is a force-to-force transfer function, it will be referred to as the *Forward Force Transfer Function*, or FFTF.

The complementary function for the return loop is then defined as $H(s) \equiv -\frac{u(s)}{F(s)}$, which includes the controller dynamics. This will be referred to as the *Return Force Transfer Function*, or RFTF. The open-loop transfer function is the product $G(s)H(s)$.

The closed-loop transfer function from command to measured force is:

$$\frac{F(s)}{r(s)} = \frac{R(s)G(s)}{1 + G(s)H(s)} \quad (3.3.2)$$

The closed-loop transfer function from command input to joint position is:

$$\frac{x(s)}{r(s)} = \frac{R(s)G(s)}{E(s)[1 + G(s)H(s)]} \quad (3.3.3)$$

where $E(s) \equiv m_P s^2 + d_E s + k_E$, the environment and payload dynamics. Actuator-axis disturbance rejection can be evaluated by looking at the closed-loop disturbance transfer function:

$$\frac{F(s)}{D(s)} = \frac{G(s)}{1 + G(s)H(s)} \quad (3.3.4)$$

3.3.2 FFTF Evaluation for this Manipulator Model

Before deriving the FFTF, some underlying assumptions will be stated. First, it is assumed for this linear analysis that the manipulator is in continuous contact with the environment. Discontinuous contact will be treated by simulation in Chapters 4 and 5. Second, it is assumed for the initial survey of control algorithms that there is no modeling error--that a controller which uses a physical parameter of the manipulator is provided with an accurate value. Note that this does not extend to the environment; no knowledge of environmental properties is incorporated into any controller studied here. One consequence of the accurate-model assumption is that the effects of viscous friction (the $d_A s$ block in Figure 3.2) can be eliminated completely with actuator velocity feedback. The inclusion of this block and its corresponding compensation becomes an unnecessary complication, and they are deleted from this part of the analysis.

Note also that only linear effects are included, so dry friction is not explicitly modeled here. The magnitude of its effect can be deduced from the disturbance rejection transfer functions; a more detailed study is performed with the simulations of Chapters 4 and 5.

Under these assumptions, equations (2.6.1) and some algebra give the following FFTF:

$$G(s) = \frac{(m_P s^2 + d_E s + k_E)(d_D s + k_D)}{(m_P s^2 + d_E s + k_E)(m_A s^2 + d_D s + k_D) + s^2[m_L m_A s^2 + (m_L + m_A)(d_D s + k_D)]} \quad (3.3.5)$$

For comparison, the FFTF for a manipulator with a rigid drive would be:

$$G(s) = \frac{m_P s^2 + d_E s + k_E}{(m_P + m_L + m_A) s^2 + d_E s + k_E} \quad (3.3.6)$$

To illustrate these functions, an example is used. The numerical parameters are chosen to reflect a relatively small manipulator with a high gear ratio (hence compliant drive) in contact with a stiff environment, carrying little or no payload. This forms a good basis for comparison because many researchers have found that stability with a rigid springlike environment is more difficult to achieve than with a primarily inertial environment^[2,5,34,54]. These values (tabulated below) will be used for a number of comparisons. Also included in the table are the frequencies of several resonances which occur in the FFTF; these are discussed below.

Parameter	Value
m_A	1 kg
m_L	1 kg
m_P	.01 kg
d_D	5 Ns/m
d_E	2 Ns/m
k_D	1000 N/m
k_E	20000 N/m
$\omega_{PLA} \equiv \sqrt{\frac{k_E}{m_P + m_L + m_A}}$	16 Hz
$\omega_P \equiv \sqrt{\frac{k_E}{m_P}}$	225 Hz
$\omega_D \equiv \sqrt{\frac{k_D}{m_A}}$	5 Hz
$\omega_E \equiv \sqrt{\frac{k_E}{m_L + m_P}}$	22.4 Hz

Table 3.1: Control Analysis Parameters

Figures 3.4 and 3.5 compare the FFTF for the above example to that of the rigid-drive case.

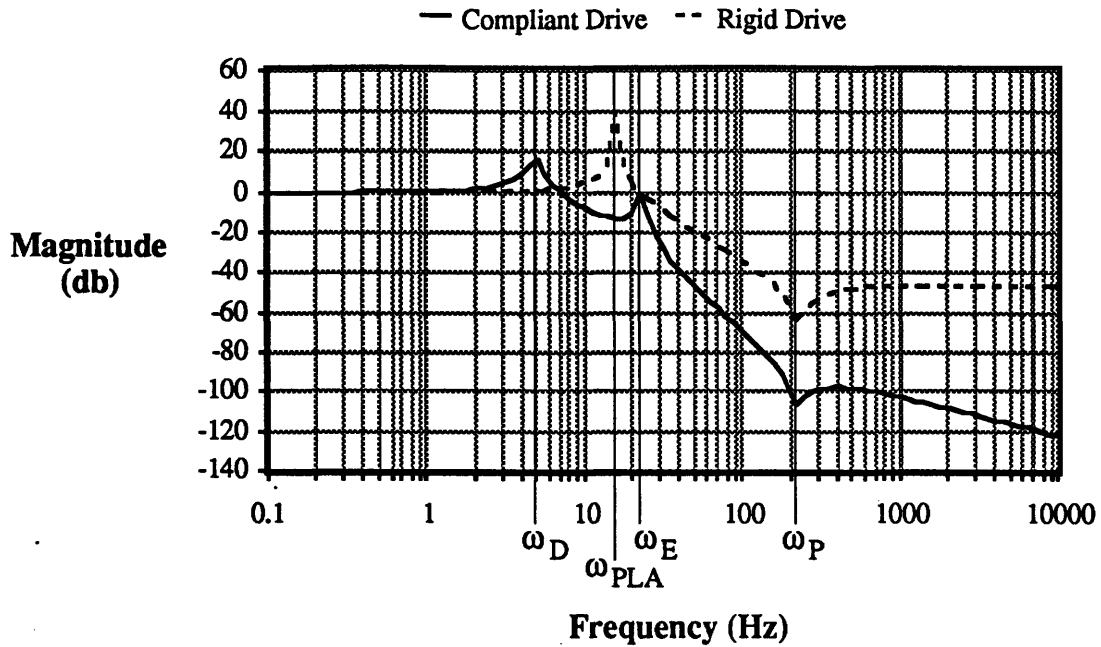


Figure 3.4: FFTF Magnitude Plot

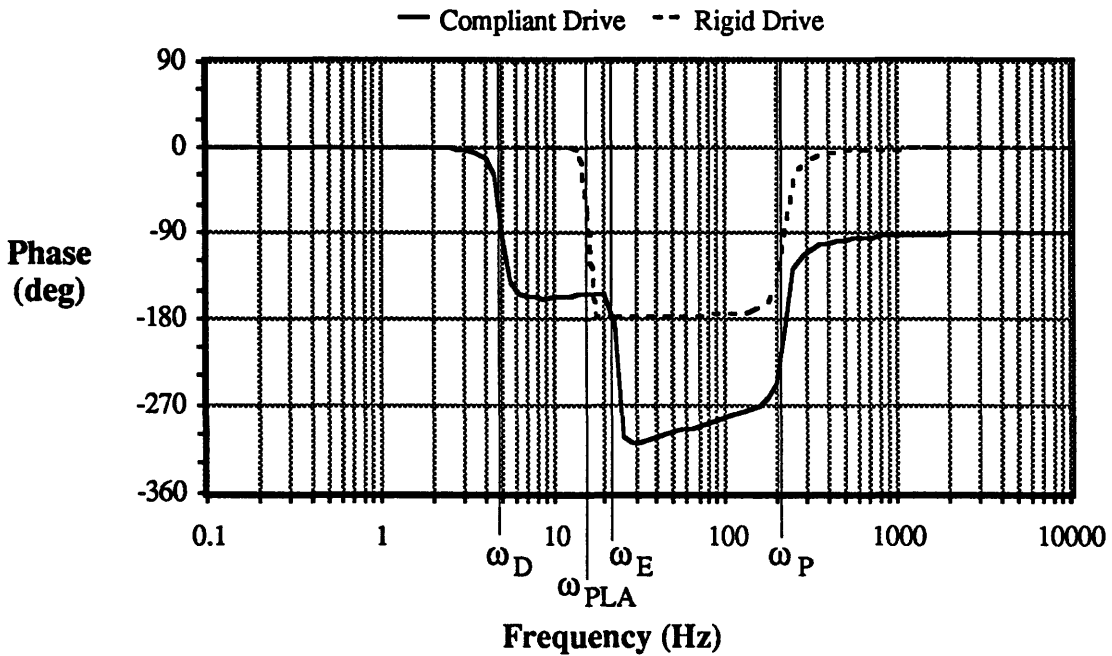


Figure 3.5: FFTF Phase Plot

Note that in both cases, $G(s) \rightarrow 1$ as $s \rightarrow 0$, as expected. In the rigid-drive case, there is a lightly damped magnitude peak at $\omega_{PLA} \equiv \sqrt{\frac{k_E}{m_P + m_L + m_A}} \approx 16$ Hz, accompanied by a 180° phase lag, and then a -40 db/decade rolloff until a magnitude dip and 180° phase lead at $\omega_P \equiv \sqrt{\frac{k_E}{m_P}} \approx 225$ Hz.

This behavior can be explained physically as follows. Since the drive is rigid, the manipulator is essentially a single mass ($m_P + m_L + m_A$) connected to a spring k_E . The expected resonance in position response X to input force U therefore occurs at the natural frequency ω_{PLA} . The FFTF, however, is the transfer function from input force U to measured force F , not to position X . The force F is related to position X through another mass-spring system, with mass m_P and spring k_E . This is an inverse resonance, because the input is position and the output is force. It occurs at frequency ω_P and appears as a trough (rather than peak) in the magnitude response, and lead (rather than lag) in the phase response. The total FFTF is then the product of the U-to-X transfer function and the X-to-F transfer function, with one resonance at ω_{PLA} and an inverse resonance at ω_P .

The response in the compliant-drive case can be understood by simplifying the FFTF in eqn. (3.3.5) with the assumption $d_D = d_E = 0$, approached in many situations of interest:

$$G(s) \approx \frac{k_D(m_P s^2 + k_E)}{m_A(m_L + m_P)s^4 + [(m_A + m_L + m_P)k_D + m_A k_E]s^2 + k_E k_D} \quad (3.3.7)$$

Defining $\omega_D \equiv \sqrt{\frac{k_D}{m_A}}$ and $\omega_E \equiv \sqrt{\frac{k_E}{m_L + m_P}}$, and substituting $i\omega$ for s :

$$G(\omega) = \left(\frac{\omega_E^2 \omega_D^2}{\omega_P^2} \right) \left\{ \frac{\omega_P^2 - \omega^2}{\omega^4 - \left[\omega_E^2 + \left(\frac{m_A + m_L + m_P}{m_L + m_P} \right) \omega_D^2 \right] \omega^2 + \omega_E^2 \omega_D^2} \right\} \quad (3.3.8)$$

In the case of an impedance-matched transmission ($m_A = m_L$) and small payload m_P , this becomes:

$$G(\omega) = \left(\frac{\omega_E^2 \omega_D^2}{\omega_P^2} \right) \left[\frac{\omega_P^2 - \omega^2}{\omega^4 - (\omega_E^2 + 2\omega_D^2)\omega^2 + \omega_E^2 \omega_D^2} \right] \quad (3.3.9)$$

From the denominator of this equation one would expect to see two resonances, each with 180° of phase lag, at frequencies $\omega^2 = \omega_D^2 + \frac{1}{2} \left(\omega_E^2 \pm \sqrt{4\omega_D^4 + \omega_E^4} \right)$. If $\omega_E \gg \omega_D$ these peaks occur at approximately ω_D and ω_E . These peaks are indeed seen in Figures 3.4 and 3.5 at $\omega_D \approx 5$ Hz and $\omega_E \approx 22.4$ Hz. The resonance trough observed at ω_P is also expected from the factor in the numerator, as in the rigid-drive case. Some additional behavior seen in the figures results from dynamics neglected in eqn. (3.3.7) and consists of a gradual 90° phase lead due to nonzero d_D , with a center frequency of $\omega = \frac{k_D}{d_D}$

≈ 32 Hz, and the corresponding magnitude slope behavior expected from a minimum-phase system.

Physically, when the drive spring k_D is much softer than the environment spring k_E , the joint motion in response to an actuator force is much smaller than actuator motion. Therefore the actuator behaves like a single mass-spring system with mass m_A and spring k_D . In the transfer function from actuator force U to actuator position Y there is a resonance at ω_D .

Actuator position essentially gives the force on the joint (again, since joint motion is small compared to actuator motion), and the joint behaves approximately like a mass-spring system with mass $(m_L + m_P)$ and spring k_E , so the transfer function from actuator position Y to joint position X has a resonance at ω_E . Combining these with the same inverse resonance at ω_P that occurs in the transfer function from X to F in the rigid-drive case gives a picture of the overall compliant-drive FFTF.

3.3.3 Consequences for Stability

The FFTF is one component of the system's open-loop transfer function. The phase shift and gain of the FFTF is added to that of the RFTF to give the whole transfer function and determine stability. The concepts of gain and phase margin can be applied to the FFTF alone to give an indication of how difficult it will be to stabilize the system with a controller, and to discover the frequencies at which the stability margin is slimmest.

The effects of drive compliance can be seen as follows. In the rigid-drive case the phase lag never exceeds -180° , so the system would theoretically be stable with a simple proportional controller with constant gain of any magnitude. If the controller itself contributes phase lag, the frequency region around ω_{PLA} is crucial, since the FFTF phase shift is near -180° and FFTF gain is around 40 db there due to the environmental resonance. The controller's RFTF would thus need to have virtually no phase lag at that frequency or at least 40 db of attenuation.

When drive compliance is present, the environmental peak is split into two: an actuator-drive resonance at ω_D and a link-environment resonance at ω_E . In the vicinity of the actuator-drive peak, the FFTF phase margin reaches a low of about 25° , constraining the phase lag (or gain) of a controller at that frequency. The most crucial frequency region is near the link-environment peak at ω_E . This is where FFTF phase lag exceeds -180° and the gain is near 0 db.

The gain margin of the FFTF (in this case 0.33 db) limits controller gains. A controller simply feeding back the measured force F multiplied by a constant gain C

(i.e. $U = -CF$; the RFTF is $H(s) = C$) would theoretically be stable for gains less than 0.33 db ($C < 1.04$). In practice, a common rule-of-thumb^[38] for controller design is to aim for a gain margin of 6 db or more (requiring $C < 0.52$), giving some stability robustness to cope with modeling errors.

In the more complicated controllers discussed in the next section, RFTF gain varies with frequency. Since the FFTF resonance peak at ω_E sets the stability limit for constant-gain controllers, it might be expected that the RFTF gain and phase behavior of other controllers specifically around this frequency would determine their stability as well. In fact, this turns out to be true for many of the controllers studied here, simplifying the stability analysis. It is relatively easy to determine the effects of parameter changes on stability when the important behavior is confined to a narrow frequency region around ω_E .

From the FFTF plot in Figure 3.4, it might appear that a stiffer environment would cause the ω_E peak to occur at a higher frequency, and the rolloff from the first peak (at ω_D) would reduce its magnitude to a less troublesome level, allowing a controller with higher gain to be stable. To investigate this Figure 3.6 was made assuming an environment ten times stiffer:




Figure 3.6: FFTF Magnitude; Stiff Environment

As the figure illustrates, the link-environment frequency ω_E increased as expected but the magnitude of the peak gain was unchanged. The gain margin of the FFTF is not improved by increasing environmental stiffness.

3.4 Force Controller Designs

This section compares several different controller designs for a given manipulator and environment to reveal their limitations, understand the tradeoffs involved, and set a performance baseline to judge improvements. No attempt is made to analyze every variation of controller design that has been proposed in the literature. The strategy is rather to proceed along a logical path from the simplest designs to ones representing the state-of-the-art. The path chosen is stiffness control, an approach with generality and an extensive literature background. This choice also reflects the philosophy of the author, who prefers to think of force control in terms of perturbations to a position trajectory.

This analysis is linear and assumes continuous environmental contact. Performance criteria are addressed as follows: stability is assessed through the constraints it imposes on gains and commanded behavior; usable bandwidth is defined, using the closed-loop transfer function, as the frequency region over which the manipulator behavior approximates the form of the commanded behavior; smoothness is compared using the closed-loop transfer function from actuator-axis disturbance to measured interaction force.

One guideline that will be used is related to robustness. In practice, the simple model used would be accurate only up to some frequency, optimistically in the 100 Hz region. Unmodeled effects such as actuator inductance, force sensor dynamics, etc. will add extra phase lag and attenuation above this frequency. If implemented digitally, the controller itself will have a time delay equal to one sampling period at best, giving phase lag proportional to frequency. So, although the FFTF shown in Figures 3.4 and 3.5 rises above -180° phase at ω_p this may not occur in the real system, and it would be unwise to base a controller on this assumption. In selecting controller gains, a safe method is to restrict the region with open-loop gain > 1 to frequencies at which the model is considered accurate, in this case 100 Hz. This rule eliminates some theoretically stable controllers which would fail in practice.

3.4.1 The Return Force Transfer Function

The RFTF is defined as $H(s) = -\frac{u(s)}{F(s)}$, with $u(s)$ given by eqn. (3.3.1). For the compliant-drive manipulator model, some algebra along with equations (2.6.1) gives:

$$H(s) = - \left[A(s) + \frac{1}{m_P s^2 + d_E s + k_E} B(s) + \frac{(m_P + m_L)s^2 + (d_E + d_D)s + k_E + k_D}{(m_P s^2 + d_E s + k_E)(d_D s + k_D)} C(s) \right] \quad (3.4.1)$$

The control law is represented by the transfer functions $A(s)$, $B(s)$, and $C(s)$. These, as shown in Figure 3.3 and eqn. (3.3.1), are the gains and compensations which the controller applies to the measured force F , joint position X , and actuator position Y , respectively. The factors which multiply them above are the transfer functions from force F to the appropriate sensor measurements.

These factors depend on the drive and environment parameters, and represent the “effectiveness” on the RFTF of the control from each type of sensor. For instance, it is evident that the effect of control based on joint position is attenuated at high frequencies, since the coefficient of $B(s)$ goes to zero as $|s|$ increases.

In each of the following sections, controller analysis begins with specifying the control law and evaluating the corresponding RFTF (eqn. 3.4.1). Comparison with the compliant-drive FFTF described in Section 3.3.2 reveals stability limits and suggests tuning rules for the gains. Choosing a stable set of gains, the performance of the controller is then investigated with the closed-loop input-output and disturbance transfer functions.

3.4.2 Explicit Force Control

One of the simplest forms of force control is explicit force control^[54]. The control law used as an example here is slightly more complicated than constant-gain proportional force feedback, but still straightforward:

$$u = F + \left[g_P + \frac{g_I}{s} \right] (F_{ref} - F) \quad (3.4.2)$$

In this implementation, the controller accepts force commands F_{ref} , and the actuator input consists of a force feedforward term (to cancel the measured force) and a

compensator applied to the force error to provide the desired motion response. A simple PI compensator is assumed. The effect of the g_P term is to give acceleration proportional to force error, and the g_I term gives acceleration proportional to integrated force error. Integral terms such as the latter are frequently used to reduce steady-state error.

The intention of an explicit force controller is not to provide a given stiffness or damping; the gains are selected to give fast stable response to force commands. The steady-state admittance ($-\frac{x(s)}{F(s)}$ as $s \rightarrow 0$) of the closed-loop system should be infinite.

Controllers of this simple type are commonly used in the literature to demonstrate theoretical points, and in some laboratory experiments^[40,41]. Hence, this is a good starting point for force controller analysis.

A block diagram for this controller is:

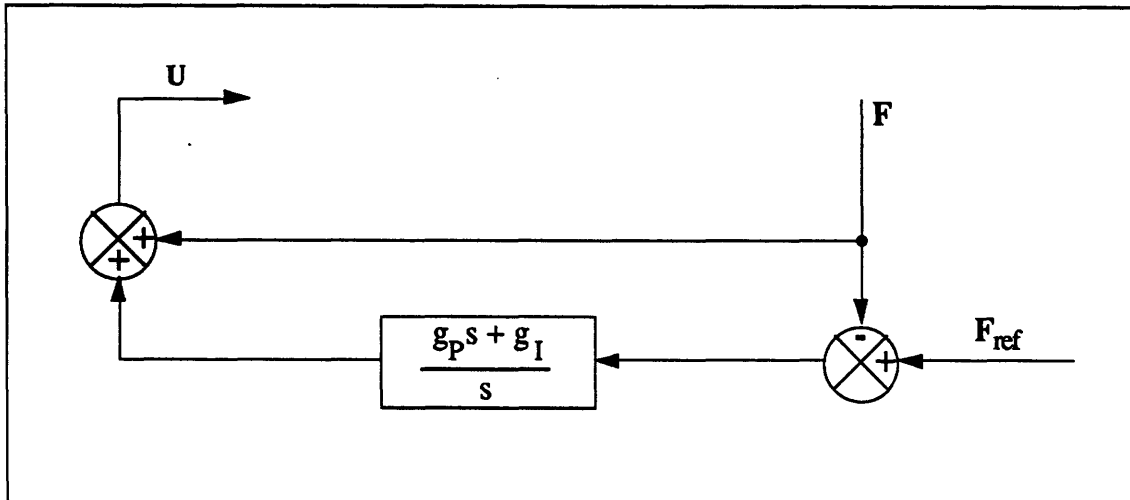


Figure 3.7: Explicit Force Controller; Block Diagram

Referring to eqn. (3.3.1) gives the following transfer functions:

$$\begin{aligned}
 A(s) &= 1 - \left[g_P + \frac{g_I}{s} \right] \\
 B(s) &= C(s) = 0 \\
 R(s) &= g_P + \frac{g_I}{s}
 \end{aligned}
 \tag{3.4.3}$$

Eqn. (3.4.1) then gives the RFTF for this controller:

$$H(s) = \frac{(g_P - 1)s + g_I}{s}
 \tag{3.4.4}$$

The following figures show the Bode plot for this RFTF:




Figure 3.8: Explicit Force Controller; RFTF Magnitude




Figure 3.9: Explicit Force Controller; RFTF Phase

The resulting open-loop transfer function (product of the FFTF and RFTF) is shown in the following figures:




Figure 3.10: Explicit Force Controller; Open-Loop Magnitude




Figure 3.11: Explicit Force Controller; Open-Loop Phase

The controller adds no phase lead, so there will still be a frequency region in the open-loop transfer function with phase $< -180^\circ$. The gains must be selected to keep the gain crossover frequency ω_{GC} out of that region. This can theoretically be done in two

ways: use high enough gains that the gain crossover occurs above the payload frequency ω_P , or use low enough gains that it occurs below the environment frequency ω_E .

The first approach leads to very high gains; $g_P \approx 200,000$. The magnitude of the open-loop transfer function is then greater than 0 db at frequencies above 200 Hz, violating the robustness constraint adopted at the beginning of this section. For this reason, the high-gain approach is rejected.

To set the gain crossover frequency below the region of excessive phase lag, note the effects of the two gains on the RFTF. The proportional gain g_P gives a constant magnitude ($= g_P - 1$), with no phase shift. The integral gain g_I contributes magnitude and 90° phase lag at frequencies below $\omega = \frac{g_I}{g_P}$. To maximize gains while retaining stability:

- Set g_P so that the ultimate gain ($g_P - 1$) of the RFTF does not exceed the gain margin of the FFTF, which occurs at ω_E and is about 0.33 db in this case. This sets the peak at ω_E in Figure 3.10 just below the unity-gain line.
- Set g_I so that phase lag of the RFTF above ω_D is less than the FFTF phase margin there, which is about 25° in this case.

These tuning rules give $g_P = 2.0$ and $g_I = 18$ in our example. For these maximum gains, the gain crossover frequency ω_{GC} is slightly higher than ω_D (7 Hz and 5 Hz, respectively, in this example). If desired, phase margin could be considerably increased by keeping ω_{GC} below ω_D . The RFTF of Figures 3.8 and 3.9, and the open-loop transfer function of Figures 3.10 and 3.11 were plotted using the maximum gains and represent performance just short of instability.

Note that in the rigid-drive case, the proportional gain would be constrained only by the robustness requirement of open-loop attenuation at 100 Hz. The integral gain is then set by the requirement to add very little phase lag at 100 Hz. This would allow the much higher gains of $g_P = 40$ and $g_I = 1000$.

The controller bandwidth attainable is shown by plotting the magnitude of the closed-loop transfer function from commanded force to measured force in Figure 3.12, using eqn. (3.3.2).




Figure 3.12: Explicit Force Control; Closed-Loop Bandwidth

It can be seen that the usable command-following bandwidth is restricted to 3 or 4 Hz (below ω_D) in this case, to avoid exciting a lightly-damped resonance in the closed-loop system at about 7 Hz (the gain crossover frequency).

This resonance peak is a consequence of the fact that these gains are close to the limit of stability. Phase margin is very small, so phase shift at the gain crossover frequency is near -180° . By definition, gain at the gain crossover frequency is unity, so the open-loop transfer function $G(s)H(s) \approx -1$ there. The denominator of the closed-loop transfer function plotted above is $1 + G(s)H(s)$. This therefore approaches 0 at ω_{GC} , giving rise to the sharp peak in response.

If gains are lowered this peak diminishes, but the closed-loop bandwidth rolls off at a lower frequency; no improvement in usable bandwidth is realized.

Disturbance rejection is shown by the magnitude of the disturbance transfer function, from eqn. (3.3.4):




Figure 3.13: Explicit Force Control; Force Disturbance Rejection

This is the magnitude of the transfer function from actuator-axis disturbance forces to output force, and should be minimized.

The resonance peak at ω_{GC} appears again here, since all of the closed-loop transfer functions share the relevant factor in the denominator. In this case, it actually causes actuator-axis disturbances at that frequency to be amplified in the output (gain > 0 db). As the next plot shows, however, commands at that frequency are amplified as well, so are not dominated by disturbances. Still, in practice it is preferable to lower the gains until this amplification of disturbances does not occur, sacrificing some disturbance rejection at lower frequencies. Overall, the disturbance rejection with this controller is minimal.

The command input, disturbance, and output are all forces, so the previous two figures can be combined to give a “signal-to-noise ratio” plot. This gives Figure 3.14, which shows the differential amplification of commands over disturbances as a function of frequency. In this plot high values are desirable, indicating that the output is influenced more by the command than the disturbance.




Figure 3.14: Explicit Force Control; S/N Ratio

This figure shows that disturbance rejection improves below 3 or 4 Hz (due to the action of the integral gain), but is not good until the frequency is quite low. If significant friction were present, requiring at least 90 % (20 db) rejection, bandwidth would have to be restricted to 0.3 Hz.

In summary, drive compliance considerably reduces the allowable gains in explicit force control. Bandwidth is limited to below ω_D , the drive resonance frequency. Disturbance rejection is poor.

3.4.3 Stiffness Control

An alternative form of force control is called *stiffness* (which here includes *damping*) control. This approach recognizes that it is more natural for many tasks to specify a position trajectory rather than a force trajectory. Interaction forces are limited to desired levels by specifying a stiffness around the reference trajectory. Multi-DOF implementations usually allow different stiffnesses in orthogonal directions. An extreme example is the hybrid control of Raibert and Craig^[40], which uses (stiff) position control in some directions and force control in others.

One practical reason for using stiffness control rather than explicit force control is the ease of including an inner position-based control loop (to improve disturbance rejection), since interaction forces are already interpreted in terms of their effects on the desired position trajectory.

The user of a stiffness-controlled manipulator needs to have a feeling for what represents a good range of stiffness values for performing tasks. In an attempt to give intuition on the meaning of various stiffness numbers, Table 3.2 presents the values for the bending stiffness of “fingers” composed of some common materials.

The finger in each case is 8 cm long, 1 cm tall, and 0.5 cm thick. If it is clamped at the base and a force applied (in the *thickness* direction) at the tip, the finger will deflect. The stiffnesses tabulated here are valid for small deflections. For comparison purposes, the approximate value for a human finger is included.

Material	Stiffness
Steel	120000 N/m
Aluminum	42000 N/m
Wood (American Elm)	5600 N/m
Soft Plastic (Nylon 11)	760 N/m
Human Finger	400 N/m

Table 3.2: Sample Stiffness Table

The simplest implementation of stiffness control (with no inner loop) is described first. The control law is:

$$u = F + (m_L + m_A) s^2 \left[x_{\text{ref}} - \frac{F}{d_T s + k_T} \right] \quad (3.4.5)$$

The actuator input consists of a force feedforward term and a term to accelerate the link to follow the desired trajectory, which is the reference modified by the measured interaction forces. No position feedback is used in this implementation (it is added in the next controller considered), so the position trajectory tracking is open-loop. The commanded stiffness and damping are k_T and d_T , respectively. Low values of these parameters are desirable, as the goal of force control is to make the manipulator more sensitive to contact forces. A block diagram for this controller is:

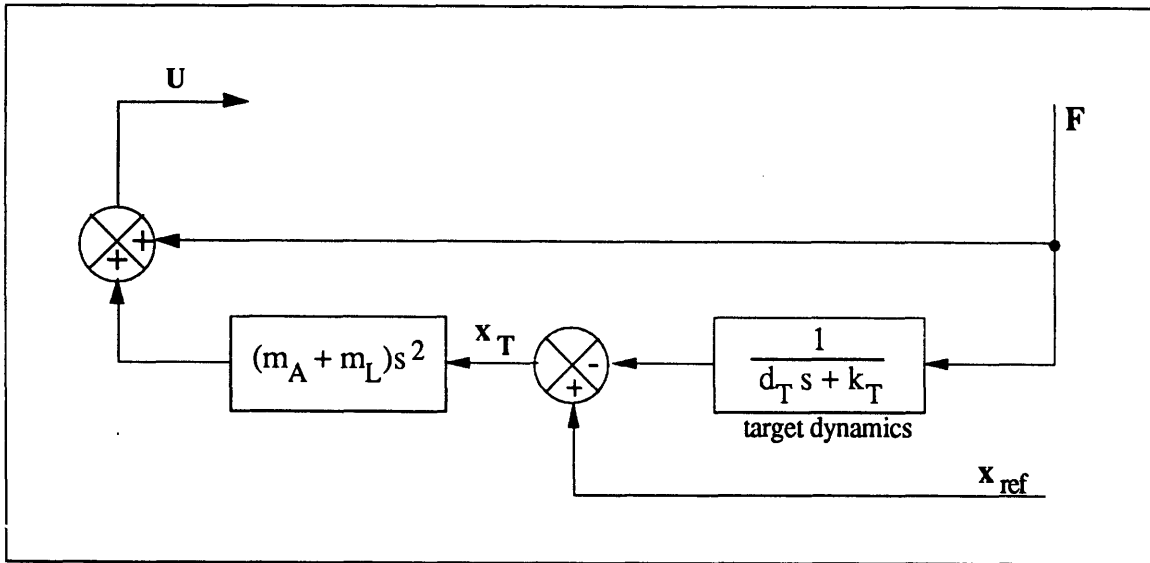


Figure 3.15: Stiffness Controller; Block Diagram

Referring to eqn. (3.3.1) gives the following transfer functions:

$$A(s) = 1 - \frac{(m_L + m_A)s^2}{d_T s + k_T}$$

$$B(s) = C(s) = 0 \quad (3.4.6)$$

$$R(s) = (m_L + m_A)s^2$$

Eqn. (3.4.1) then gives the RFTF for this controller:

$$H(s) = \frac{(m_L + m_A)s^2 - d_T s - k_T}{d_T s + k_T} \quad (3.4.7)$$

Figures 3.16 and 3.17 show the Bode plot for this RFTF:




Figure 3.16: Stiffness Controller, RFTF Magnitude




Figure 3.17: Stiffness Controller, RFTF Phase

The resulting open-loop transfer function (product of the FFTF and RFTF) is shown in the following figures:




Figure 3.18: Stiffness Controller; Open-Loop Magnitude




Figure 3.19: Stiffness Controller; Open-Loop Phase

The crucial frequency is ω_E , near which the FFTF phase falls to -310° and magnitude is approximately 0 db. Two cases emerge in the RFTF, depending on whether $k_T(m_L + m_A) > d_T^2$ or not. If so the desired behavior is “springlike,” and the magnitude of the RFTF at ω_E is inversely proportional to k_T if $\frac{k_T}{d_T} > \omega_E$, or inversely proportional to

d_T otherwise. The RFTF phase lead at ω_E increases with $\frac{k_T}{d_T}$. If $k_T(m_L + m_A) < d_T^2$ the desired behavior is “damperlike,” and the magnitude of the RFTF at ω_E is inversely proportional to d_T . The RFTF phase lead increases with d_T .

This enforces a tradeoff between k_T and d_T . For damperlike behavior, a low d_T requires a high k_T and vice-versa. In the present case, $d_T > 140$ Ns/m allows arbitrarily low k_T , and $k_T > 5000$ N/m allows arbitrarily low d_T . The RFTF plotted above used intermediate values of $d_T = 120$ Ns/m and $k_T = 3000$ N/m. For very springlike behavior, stability can obtain despite low values of k_T and d_T , but only if the ratio $\frac{k_T}{d_T}$ is undesirably large. Figure 3.20 illustrates this tradeoff:




Figure 3.20: Stiffness Controller; Stability Tradeoff

The dotted line in the figure shows combinations of k_T and d_T which give a constant damping ratio, in this case selected to give critical damping to a payload of 1 kg. It is included as a reminder that the desired behavior may be a pure damper, but not an underdamped spring (oscillation is to be avoided). Hence the range of useful commanded behavior for this controller is the shaded region, described as stable and overdamped. Points near the upper edge of this region give springlike behavior, and points near the bottom give damperlike behavior.

Note that a more compliant environment would improve stability by decreasing ω_E , which places it in a region of greater RFTF phase lead and decreased RFTF gain. This effect is illustrated in Figure 3.21, which shows the minimum target damping d_T which can be stably achieved as a function of environmental stiffness k_E (assuming $k_T = 0$):

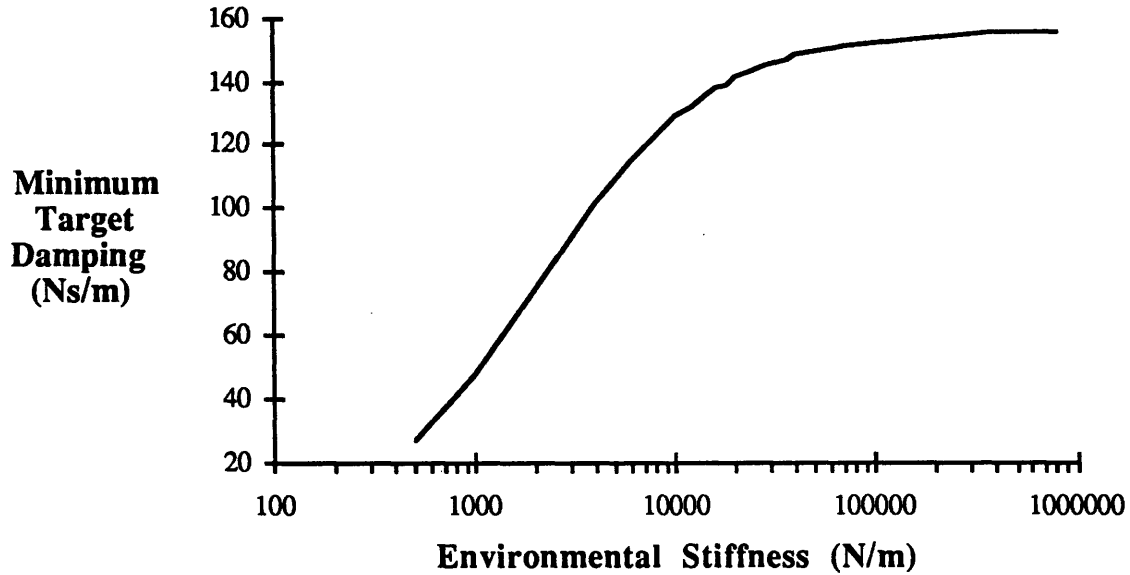


Figure 3.21: Stiffness Controller, Target Damping Tradeoff

To show closed-loop performance of the stiffness controller, the closed-loop admittance $I(s)$ is plotted in Figure 3.22. This is defined as $-\frac{x(s)}{F(s)}$ with the environment removed from the loop, and represents the position response to an applied tip-force:

$$I(s) \equiv -\frac{x(s)}{F(s)} = \quad (3.4.8)$$

$$\frac{m_A s^2 + d_D s + k_D - C(s) - (d_D s + k_D)A(s)}{(m_L s^2 + d_D s + k_D)[m_A s^2 + d_D s + k_D - C(s)] - (d_D s + k_D)[B(s) + d_D s + k_D]}$$

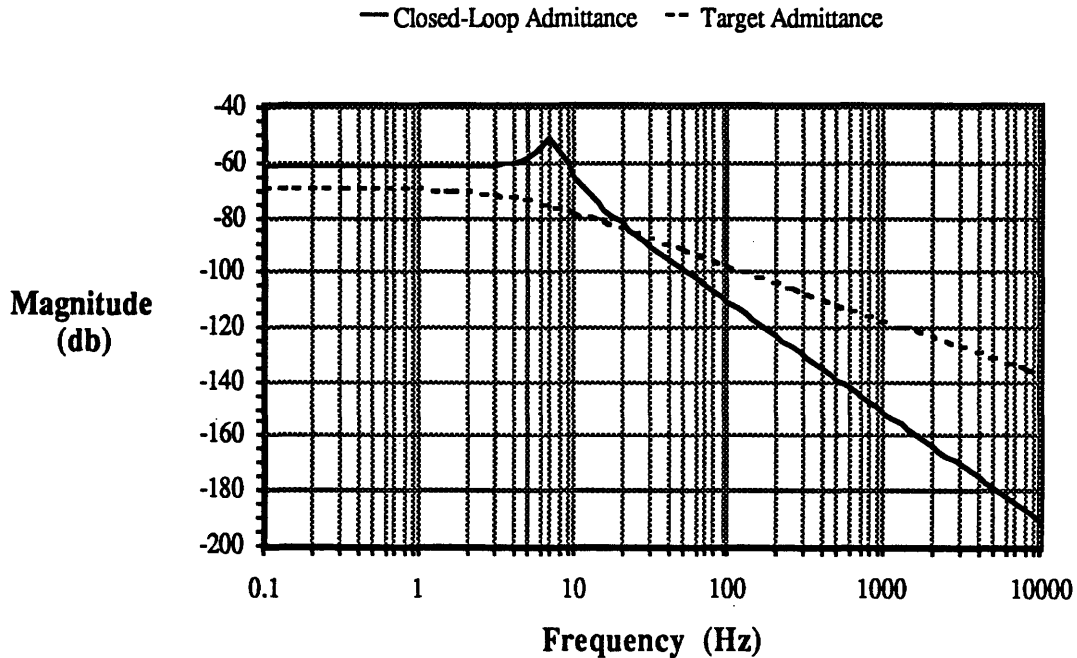


Figure 3.22: Stiffness Controller; Closed-Loop Admittance

The target admittance $\frac{1}{d_T s + k_T}$ is plotted also for reference. Two things are clear from this figure: usable command bandwidth is again limited to about 4 Hz, below the drive resonance ω_D , and there is an error in the closed-loop stiffness at low frequencies. This is due to the (unmodeled) drive compliance, in series with the active compliance. In this case the target steady-state stiffness $k_T = 3000$ N/m (corresponding to -69.5 db on the chart), yet the achieved steady-state value is around 1200 N/m (-62 db).

The magnitude of the disturbance transfer function is shown in Figure 3.23. It can be seen that the control law amplifies low-frequency actuator disturbances. The reason is that the force-feedback term in the control law is essentially a high-pass filter, and at low frequencies its effect is dominated by the force feedforward term. The feedforward term could be deleted, giving a unity-gain disturbance transfer function at low frequencies, but the closed-loop admittance would then fail to approximate the target admittance at low frequencies.




Figure 3.23: Stiffness Controller; Force Disturbance Rejection

In summary, the simple stiffness controller is also limited in bandwidth to below the drive resonance frequency ω_D . Accuracy of closed-loop admittance suffers when commanded to exceed drive stiffness, although “damperlike” behavior can be accurately produced. The stability limits enforce a tradeoff between d_T and k_T , restricting commanded behavior to be rather unyielding either as a damper or a spring. Environmental properties affect stability, requiring stiffer commanded dynamics with a stiffer environment. Disturbance rejection is practically nonexistent.

These factors combine to defeat the utility of the force feedback loop. It is required to be stiff for stability, yet in that case drive compliance dominates the “active” compliance. With no position feedback, position tracking and disturbance rejection is poor. Performance is basically that of an open-loop position controller, with whatever passive compliance is provided by the transmission.

3.4.4 Stiffness Control with Inner Loop

The limitations of a simple stiffness controller quickly became apparent to researchers trying to implement force control, sparking much of the recent research on dynamic modeling and design improvements^[34,54]. The first major improvement, adopted widely^[34,47,51], is to add a position- (or velocity-) based inner control loop. The idea is that a high-gain colocated-sensor inner loop will improve disturbance rejection, while hopefully not degrading stability (since it does not enclose the unmodeled drive compliance or environment, etc.). There is the added advantage that many commercially available manipulators already have such loops at each joint, and can be converted to force controllers relatively easily.

The control law is:

$$u = F + (g_D s + g_P) \left[x_{\text{ref}} - \frac{F}{d_T s + k_T} - y \right] \quad (3.4.9)$$

Actuator input consists of the usual force feedforward term and a term to drive y (the actuator position) to match the desired trajectory, which is again the reference modified by measured interaction forces. The parameters g_D and g_P are derivative and proportional gains of the position-based inner loop, respectively. Other types of inner-loop compensation are used, often with an integrator to eliminate static disturbances, but for this analysis a PD compensator was considered sufficiently general. Acceleration feedforward in the form of the second term of eqn. (3.4.5) is not included because it requires additional differentiation of the force measurement, adding noise and unwanted dynamics. Reliance is made instead on inner-loop feedback to track the position trajectory.

A block diagram for this controller is:

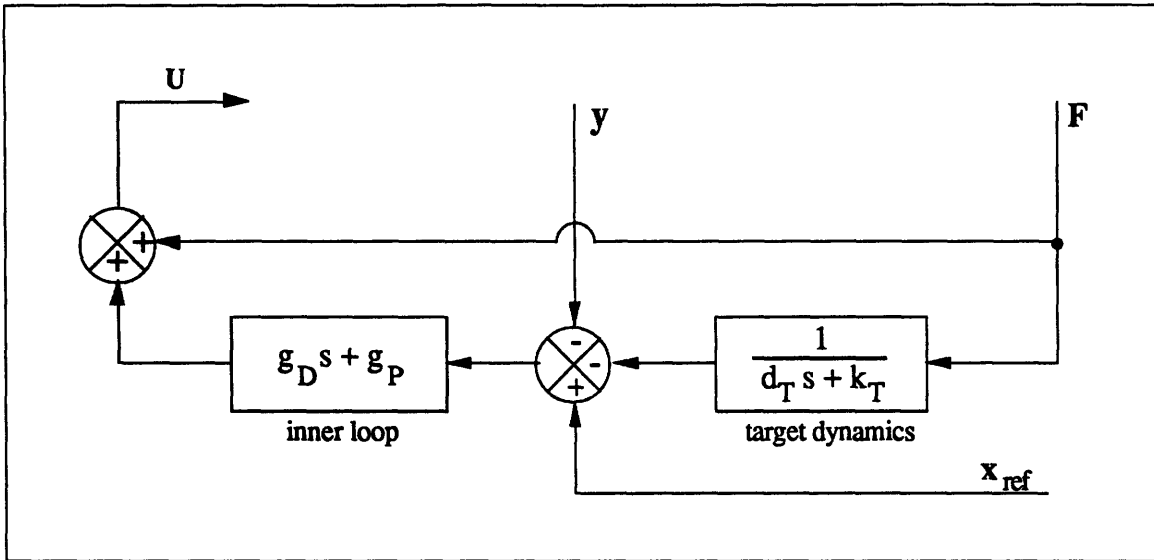


Figure 3.24: Stiffness Controller w/Inner Loop; Block Diagram

Referring to eqn. (3.3.1) gives the following transfer functions:

$$A(s) = 1 - \left(\frac{g_D s + g_P}{d_T s + k_T} \right)$$

$$B(s) = 0 \quad (3.4.10)$$

$$C(s) = - (g_D s + g_P)$$

$$R(s) = g_D s + g_P$$

Eqn. (3.4.1) then gives the RFTF for this controller:

$$H(s) = \left(\frac{g_D s + g_P}{d_D s + k_D} \right) \left[\frac{(m_P + m_L) s^2 + (d_E + d_D) s + k_E + k_D}{m_P s^2 + d_E s + k_E} + \frac{d_D s + k_D}{d_T s + k_T} \right] - 1$$

(3.4.11)

The following figures show the Bode plot for this RFTF:




Figure 3.25: Stiffness Controller w/Inner Loop; RFTF Magnitude




Figure 3.26: Stiffness Controller w/Inner Loop; RFTF Phase

To plot the RFTF shown above, the same target values were used as for the previous controller: $k_T = 3000$ N/m and $d_T = 120$ Ns/m. The inner-loop damping ratio was set at $\zeta_I = .707$ and the inner-loop bandwidth was increased to a point near the stability limit, giving $\omega_I = 22$ Hz.

The resulting open-loop transfer function (product of the FFTF and RFTF) is shown in the following figures:




Figure 3.27: Stiffness Controller w/Inner Loop; Open-Loop Magnitude




Figure 3.28: Stiffness Controller w/Inner Loop; Open-Loop Phase

The appearance of this RFTF is quite different from that of the stiffness controller without an inner loop (Figures 3.16 and 3.17). The dominant behavior now comes from the actuator-position term in equation (3.4.1). The coefficient of $C(s)$ in that equation is

$\left[\frac{(m_P + m_L)s^2 + (d_E + d_D)s + k_E + k_D}{(m_P s^2 + d_E s + k_E)(d_D s + k_D)} \right]$, and contributes the observed resonance at ω_P and the inverse resonance at ω_R , where ω_R is defined by:

$$\omega_R \equiv \sqrt{\frac{k_E + k_D}{m_P + m_L}} \quad (3.4.12)$$

The actuator-position term (the first term in the bracket in eqn. 3.4.11) gives approximately 180° phase lead between the frequencies ω_R and ω_P , accompanied by a sharp dip in magnitude near ω_R and a sharp peak at ω_P . This almost matches an inverse feature in the FFTF (see Figures 3.4 and 3.5). If the match were perfect, the open-loop transfer function phase would never drop below -180° and stability would be endangered only by unmodeled dynamics. Unfortunately, ω_R , the frequency at which the 180° phase lead of the RFTF occurs, is *slightly higher* than $\omega_E \equiv \sqrt{\frac{k_E}{m_P + m_L}}$, the frequency at which the FFTF phase falls 180° . The separation is due to drive compliance k_D . In this example, $\omega_E = 22.40$ Hz and $\omega_R = 22.95$ Hz.

The result is a narrow region between ω_E and ω_R where the phase of the open-loop transfer function falls abruptly and then returns. Instability can occur as follows: Just below ω_R the phase of the open-loop transfer function falls below -180° . At ω_R the RFTF magnitude drops to a small value set by the low level of damping in drive and environment, causing the open-loop transfer function magnitude to cross 0 db. When RFTF magnitude increases again on the high-frequency side of ω_R , the open-loop transfer function phase is above -180° due to the rising phase of the RFTF; the result is an encirclement of the -1 point by the Nyquist path.

In the plots shown above, this does not quite happen; the gains were chosen to give stability. The phase of the open-loop transfer function only descends to -173° . However, its gain crossover frequency is near ω_E , and the potential for this type of instability can be seen.

For this controller it is again the response near ω_E that determines stability. In root-locus terms, it is the link-environment poles which first cross the imaginary axis into the right half-plane. As a result, we can now concentrate on the vicinity of ω_E to evaluate the effects of parameter changes on stability. Changes which decrease phase lag in the RFTF at ω_E improve stability, and changes which increase it bring on instability.

The first term (representing actuator position feedback) in the RFTF of eqn. (3.5.11) is multiplied by the factor $\left[\frac{g_D s + g_P}{d_D s + k_D} \right]$. Define corner frequencies $\omega_{Ci} \equiv \frac{g_P}{g_D} = \frac{\omega_I}{2\zeta_I}$

and $\omega_{Cd} \equiv \frac{k_D}{d_D} = \frac{\omega_D}{2\zeta_D}$ where ω_I and ζ_I are the inner-loop bandwidth and damping ratio, respectively and ζ_D is the drive damping ratio ($\approx .1$ in this case). The denominator approximately gives 90° of phase lag above frequency ω_{Cd} , and the numerator gives 90° of phase lead above frequency ω_{Ci} . The net effect of this factor on phase is therefore to give 90° lag between the frequencies ω_{Cd} and ω_{Ci} if $\omega_{Ci} > \omega_{Cd}$, or 90° lead if $\omega_{Cd} > \omega_{Ci}$. Stability is thus improved if $\omega_{Cd} > \omega_E > \omega_{Ci}$, and the wider the frequency separation, the better (since the 90° phase transitions do not occur immediately at the corner frequency).

It follows that if the drive is stiffened ω_D increases, so ω_{Cd} increases, improving stability. If inner-loop bandwidth ω_I is increased ω_{Ci} increases, degrading stability. Stiffening the environment also degrades stability, since for these parameters ω_E is close to ω_{Cd} and will exceed it if slightly increased, decreasing the phase lead due to the first term in the RFTF.

The target dynamics enter into the second term of the RFTF (the term representing force feedback): $\left[\frac{g_{DS} + g_P}{d_{TS} + k_T} \right]$. Defining corner frequency $\omega_{Ct} \equiv \frac{k_T}{d_T}$, this term gives 90° phase lag between the frequencies ω_{Ct} and ω_{Ci} if $\omega_{Ci} > \omega_{Ct}$, or 90° phase lead if $\omega_{Ct} > \omega_{Ci}$. Stability is improved if $\omega_{Ct} > \omega_E > \omega_{Ci}$.

This is a constraint on commanded behavior which is difficult to meet. The inequality $\omega_{Ct} > \omega_{Ci}$ requires the ratio of k_T to d_T to be so large that the manipulator would act like a very lightly-damped spring; such oscillatory behavior is undesirable. As a result, the second term of the RFTF usually contributes lag instead of lead. This effect can be reduced by increasing $(d_T^2 \omega_E^2 + k_T^2)$, which diminishes the overall magnitude of the term. This leads again to a tradeoff between lowest achievable values of k_T and d_T , which is illustrated in Figure 3.29.

The highest inner-loop bandwidths (desirable for increased disturbance rejection) require a high target stiffness and a soft environment for stability. The first term dominates the RFTF in this case. Lowering the target values d_T or k_T always decreases the bandwidth that can be achieved, by increasing the influence of the destabilizing second term.




Figure 3.29: Stiffness Controller w/Inner Loop; Stability Tradeoff

The closed-loop admittance is plotted below:




Figure 3.30: Stiffness Controller w/Inner Loop; Closed-Loop Admittance

The admittance achieved is very similar to that of the previous controller. The steady-state closed-loop stiffness attained is simply the commanded stiffness in series with the drive stiffness:

$$k_{T,achieved} = \frac{k_D k_T}{k_D + k_T} \quad (3.5.13)$$

Attempts to command behavior stiffer than the drive are therefore fruitless. Usable bandwidth is again limited to approximately the drive frequency ω_D .

The disturbance transfer function is given by equation (3.3.4). For this controller it becomes:

$$\frac{F(s)}{D(s)} = \frac{(d_{TS} + k_T)(m_P s^2 + d_{ES} + k_E)(d_{DS} + k_D)}{(d_{TS} + k_T) \left[((m_P + m_L) s^2 + (d_E + d_D) s + k_E + k_D)(m_A s^2 + g_{DS} + g_P) + \dots \right.} \quad (3.4.14)$$

$$\left. \dots (d_{DS} + k_D) m_L s^2 \right] + (g_{DS} + g_P)(m_P s^2 + d_{ES} + k_E)(d_{DS} + k_D)$$

The magnitude of the disturbance transfer function is shown below:




Figure 3.31: Stiffness Controller w/Inner Loop; Force Disturbance Rejection

It is clear from comparison with Figure 3.23 that disturbance rejection below ω_I , the inner-loop bandwidth, is much improved, achieving in this case nearly 30 db attenuation throughout the usable bandwidth. Disturbance rejection above this frequency is essentially unchanged.

In summary, the addition of an inner position loop to the simple stiffness controller improves performance by allowing higher disturbance rejection and widening the range of stable commanded behavior. With such an inner loop, increasing the inner-loop bandwidth ω_I improves disturbance rejection, but degrades stability. However, it is possible to find gains which give much improved disturbance rejection *and* wider stability limits than the simple stiffness controller of the previous section.

Other features remain the same. Overall bandwidth is still limited to below the drive resonance ω_D , and inaccuracy in the achieved admittance occurs when commanded stiffness exceeds drive stiffness. As before, increasing drive compliance or environmental stiffness degrades stability. There is still a tradeoff between the lowest allowable values of d_T and k_T .

3.4.5 Stiffness Control with Inner Loop and Filtering

The major limitation of the controllers so far described is the inability to attain low values of k_T and d_T . Low target stiffness means the manipulator must react strongly to measured forces. This means large force feedback, which is destabilizing for this system since the force sensor is non-colocated. In the stiffness controller with inner loop, force feedback is the second term in brackets in eqn. (3.4.11). As described above, low values of k_T and d_T increase the influence of this term on the RFTF at the link-environment natural frequency ω_E , which increases phase lag and leads to instability.

One approach to this problem is filtering. If a low-pass filter is placed in series with the force feedback, the magnitude of this term at the stability-critical frequency ω_E can be reduced while retaining the same force feedback at lower frequencies.

The filter cutoff frequency is chosen to give sufficient attenuation at ω_E . With a stiff environment, a high filter cutoff could be used. However, for robustness, the same controller should be stable with softer environments. In particular, to avoid the need to switch controller gains at environmental contact it should be stable for $k_E = d_E = 0$ (free space). If ω_E decreases below the actuator-drive natural frequency ω_D , ω_D becomes the crucial frequency for stability (the frequency at which FFTF phase falls below -180°). Thus, for robustness, the filter cutoff should be below ω_D . Since the usable bandwidth of

the non-filtered controller is limited to less than ω_D anyway, tracking of the commanded behavior should not be impaired.

The control law is:

$$u = F + (g_{Di}s + g_{Pi}) \left[x_{ref} - \left(\frac{g_{Do}s + g_{Po}}{m_L s^2 + g_{Do}s + g_{Po}} \right) \left(\frac{F}{d_T s + k_T} \right) - y \right] \quad (3.4.15)$$

Here the particular form assumed for the low-pass filter is $\left(\frac{g_{Do}s + g_{Po}}{m_L s^2 + g_{Do}s + g_{Po}} \right)$ with two adjustable gains g_{Do} and g_{Po} . The use of m_L in the denominator serves to normalize the gains and is unimportant at this point; it is there to facilitate later comparisons.

Some filtering of analog signals (e.g. those from a force sensor) before sampling is commonly used in practice to prevent aliasing^[6] and reduce noise. The filter frequencies used for these purposes are relatively high. Anti-aliasing filters have a cutoff around half the sampling frequency, which is itself typically 20-30 times the bandwidth of the inner loop. The filter suggested here for control purposes has a cutoff frequency near the overall bandwidth of the controller; around a factor of 10 lower than inner-loop bandwidth and roughly 100 times lower than that of an anti-aliasing filter.

Controllers of this category have been used by several researchers, although the particular form of the filter and the reasoning cited varies. This principle helps explain the good force control performance found empirically by Maples and Becker^[34] in their experiments. Hogan's impedance control^[21,22,56] obtains the same effect by imposing a target mass on the desired behavior.

A block diagram for this controller is:

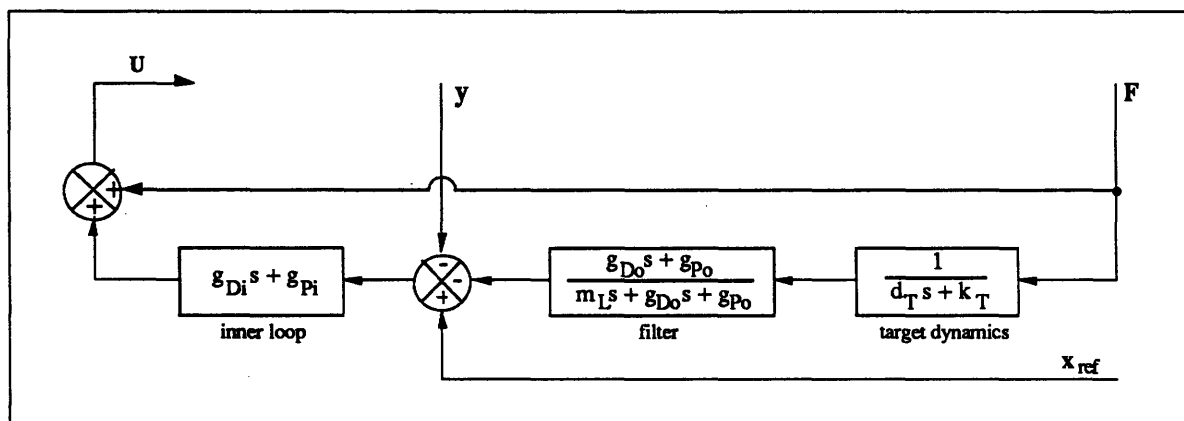


Figure 3.32: Stiffness Controller w/Inner Loop + Filter; Block Diagram

Referring to eqn. (3.3.1) gives the following transfer functions:

$$\begin{aligned}
 A(s) &= 1 - \left(\frac{g_{D0}s + g_{P0}}{m_L s^2 + g_{D0}s + g_{P0}} \right) \left(\frac{g_{Di}s + g_{Pi}}{d_T s + k_T} \right) \\
 B(s) &= 0 \\
 C(s) &= - (g_{Di}s + g_{Pi}) \\
 R(s) &= g_{Di}s + g_{Pi}
 \end{aligned} \tag{3.4.16}$$

Eqn. (3.4.1) then gives the RFTF for this controller:

$$\begin{aligned}
 H(s) &= \left(\frac{g_{Di}s + g_{Pi}}{d_D s + k_D} \right) \left[\frac{(m_P + m_L)s^2 + (d_E + d_D)s + k_E + k_D}{m_P s^2 + d_E s + k_E} \right. \\
 &\quad \left. + \left(\frac{g_{D0}s + g_{P0}}{m_L s^2 + g_{D0}s + g_{P0}} \right) \left(\frac{d_D s + k_D}{d_T s + k_T} \right) \right] - 1
 \end{aligned} \tag{3.4.17}$$

Figures 3.33 and 3.34 show the Bode plot for this RFTF. Filter bandwidth was set at 6 Hz with a damping ratio of .707. Inner-loop bandwidth was increased to 50 Hz, and the commanded stiffness was reduced to $k_T = 600$ N/m and $d_T = 60$ Ns/m. Compared with Figures 3.25 and 3.26, attenuation and phase lag in the force feedback term of the RFTF due to the filter is visible at low frequencies. Because of the attenuation the collocated position-feedback term dominates at ω_E , giving a net phase lead there over the non-filtered controller, as intended. Except for these effects, behavior is similar to the previous controller.

Figures 3.35 and 3.36 show the open-loop transfer function. The same potential for instability exists around the link-environment frequency ω_E . This leads to a stability tradeoff (illustrated in Figure 3.37) which appears qualitatively similar to that of Figure 3.29. Note, however, that the filter permits numerically lower stiffnesses to be achieved.

The closed-loop admittance is shown in Figure 3.38. Since the commanded stiffness is less than the drive stiffness in this case, the achieved steady-state value is more accurate than before. The stiffness obtained is 375 N/m for a target value of 600 N/m, as expected from equation (3.4.13). Usable bandwidth is still around 4 Hz with the filter.




Figure 3.33: Stiffness Controller w/Inner Loop + Filter; RFTF Magnitude




Figure 3.34: Stiffness Controller w/Inner Loop + Filter, RFTF Phase




Figure 3.35: Stiffness Controller w/Inner Loop + Filter; Open-Loop Magnitude




Figure 3.36: Stiffness Controller w/Inner Loop + Filter; Open-Loop Phase




Figure 3.37: Stiffness Controller w/Inner Loop + Filter; Stability Tradeoff




Figure 3.38: Stiffness Controller w/Inner Loop + Filter; Closed-Loop Admittance

The disturbance transfer function is given by equation (3.3.4). For this controller it becomes:

$$\frac{F(s)}{D(s)} = \frac{(d_{Ts} + k_T)(m_{ps}^2 + d_{Es} + k_E)(d_{Ds} + k_D)(m_L s^2 + g_{D0}s + g_{P0})}{(d_{Ts} + k_T)(m_L s^2 + g_{D0}s + g_{P0}) [((m_P + m_L)s^2 + (d_E + d_D)s + k_E + k_D) \cdot \dots \dots (m_A s^2 + g_{Di}s + g_{Pi}) + (d_{Ds} + k_D)m_L s^2] + \dots \dots (g_{D0}s + g_{P0})(g_{Di}s + g_{Pi})(m_{ps}^2 + d_{Es} + k_E)(d_{Ds} + k_D)}$$
(3.4.18)

Figure 3.39 shows the magnitude of this function. Since the inner-loop bandwidth could be increased to 50 Hz, disturbance rejection improved to about 50 db at frequencies below 1 Hz.




Figure 3.39: Stiffness Controller w/Inner Loop + Filter; Force Disturbance Rejection

In summary, the addition of filtering to the force feedback term improved stability, by reducing the non-coloated force sensor's influence on open-loop phase near the gain crossover frequency. Inner-loop bandwidth could then be increased, improving

disturbance rejection. At the same time, lower values of k_T and d_T could be attained, making the manipulator more responsive to interaction forces. The required filtering diminished the usable bandwidth only slightly, still limited to below the actuator-drive natural frequency ω_D (which is 5 Hz in this example).

3.5 Conclusions

This chapter set criteria for good performance, and studied the workings of a series of force control algorithms applied to the compliant-drive manipulator model. One feature at a time was added to the controllers, starting with the simplest and ending with one which represents the best-performing current design.

The analysis establishes that the simple manipulator model used can produce the kind of performance limitations seen in practice, and that a frequency-domain approach to controller analysis is useful in this context. It is shown that these methods can explain the successes and failures of present control algorithms, in particular clarifying how force feedback filtering improves stability.

Dividing the system into Forward and Return Force Transfer Functions is a good way to compare different controllers for the same system, develop tuning rules for gains, and discover ways to improve stability. One important result for the controllers shown is that stability limits can be understood solely by examining the behavior of the RFTF in the vicinity of ω_E , the link-environment resonance frequency. When the transfer function is evaluated at a particular frequency, the effect of each term is clear. It is then easier to predict the stability consequences of parameter changes.

To summarize the important features seen in force controller behavior: Gains are significantly limited by drive compliance. Bandwidth is limited to ω_D , the drive-actuator resonance frequency. The upper limit on achieved stiffness is k_D , the drive stiffness.

In stiffness controllers, the force feedback term contributes phase lag to the RFTF at frequencies above ω_D . More compliance in the drive or less in the environment degrades stability, by lowering ω_D in the first case or raising ω_E (which is approximately the gain crossover frequency of the open-loop transfer function) in the second. Either of these has the effect of increasing the phase lag of the RFTF (and therefore the open-loop transfer function) at ω_E . Lowering inner-loop bandwidth or increasing commanded stiffness is required to regain stability.

There is a tradeoff in lowest achievable values of the commanded parameters k_T and d_T . Increasing either parameter reduces the magnitude of the destabilizing force-feedback term.

A position-based inner control loop greatly improves disturbance rejection within its bandwidth, but high gains still lead to instability (even though the primary inner-loop sensor is colocated). Low-pass filtering of the force feedback to give attenuation at ω_E allows higher inner-loop gains and lower values of k_T and d_T . Essentially this decouples the desirable effects of force feedback at frequencies below ω_D from the undesirable effects on stability at the higher frequency ω_E .

The final controller studied in this chapter (the stiffness controller with inner loop and filtering) is taken as a reference design. It will be used to evaluate the ideas suggested in the next chapter.

Chapter 4: Joint/Actuator Controller Design

This chapter presents two interrelated ideas for improving the force-control performance of a manipulator with drive compliance. One is a change in mechanical design of the manipulator, and the other is a new control algorithm. The success of either idea requires the use of both. In each case, the key concept is the separate treatment of joint motion and actuator motion; hence the term “Joint/Actuator Controller.”

The first recommendation of this thesis is to further mechanically decouple the joint from the actuator by *adding* drive compliance. This may seem strange, as drive compliance has been cast as a problem up to this point. Section 4.1 details the expected benefits and drawbacks.

The control algorithm described in Section 4.2 relies on this mechanical decoupling and an internal model of the drive compliance, and is intended to overcome the bandwidth limit set by drive frequency ω_D . Stability range is also improved, allowing lower target stiffnesses to be achieved.

Section 4.3 compares the resulting performance of the Joint/Actuator design with the best-performing design of Chapter 3: the stiffness controller with inner loop and filtering. The comparison is made using the performance criteria defined in Section 3.1.2.

The final section of the chapter summarizes the comparison results, which show improved performance in every category.

4.1 Mechanical Decoupling of Joint and Actuator Dynamics

Decoupling here implies that drive stiffness is made low enough that actuator position and joint position must be considered distinct variables in the operating frequency range. This can be accomplished by simply adding a linear spring in series with the transmission. The joint force is then determined by actuator position relative to joint position, rather than directly by actuator force. Increasing drive compliance is intended to improve performance in the areas of smoothness and contact behavior.

The next two sections explain these benefits. The third section discusses some other effects of drive decoupling, and the fourth section describes the principal cost of this approach--increased actuator force requirements.

4.1.1 Smoothness Benefits

Smoothness is improved in several ways. A more compliant drive reduces the effect of actuator force disturbances (stiction, friction, cogging, etc.) on tip forces, when there is an inner loop based on actuator position. This works as follows: a disturbance force at the actuator gives rise to a certain position error: the magnitude of the force divided by the proportional gain of the inner loop (for steady-state error). When in contact with a rigid environment, the resulting error in force applied at the tip is the position error multiplied by the drive stiffness. With a stiff drive, small position errors can give large tip-force errors.

Of course, closing the outer loop on the tip force sensor acts to reduce this error. This is done, however, with the relatively low gains of the outer loop. By decoupling the drive, the outer loop is faced with significantly smaller force errors to correct. Hence, the rejection of actuator disturbance forces is improved.

Also, the analysis in Section 2.3 indicates that an actuator limit cycle caused by backlash can be thought of as a position disturbance rather than a force disturbance, since it approaches a minimum amplitude which is independent of inner-loop gains. The most direct way to reduce the effect of such a disturbance on tip force is to increase drive compliance.

Cogging, backlash, and the stick-slip behavior caused by friction introduce unwanted high-frequency components into actuator motion. Drive compliance improves smoothness by filtering these out of the endpoint trajectory.

With drive decoupling, the smallest forces that can be applied open-loop (without force feedback) at the manipulator tip are determined by the positioning resolution of the actuator, rather than by the magnitude of disturbance forces at the actuator. For instance, a joint force increment smaller than the level of actuator friction can be reliably applied, given a high-resolution inner loop and a compliant drive. The maximum static force that can be applied is still the same, so that force dynamic range is improved.

4.1.2 Contact Behavior Benefits

For environmental contact, consider a rigid-drive manipulator touching a rigid surface. The entire manipulator must stop immediately, generating large interaction forces. With compliant drives, less mass must be stopped by the initial contact, giving the controller more time to react and safely bring the rest of the manipulator to a halt.

The reduction in inertia depends on configuration. Links adjacent to the wrist may be aligned with the direction of motion. These contribute to wrist mass and stop immediately, since the drive compliance for these links is not deflected at contact. The link drives which are *producing* motion in the contact direction, however, always act to decouple inertias during contact, so smaller interaction forces are generated than with a rigid drive. If some drive compliance already exists, increasing it will lower forces still further.

4.1.3 Other Performance Effects

Increasing drive compliance has further consequences. If a simple linear spring is added in series with the transmission, both the drive resonance frequency ω_D and damping ratio ζ_D decrease. As seen in Chapter 3, this limits performance of the controllers discussed so far, which are implicitly based on a rigid-drive manipulator model. For a controller based on a compliant-drive model (as discussed in Section 4.2) adding a spring to the transmission can actually improve performance by enhancing modeling accuracy. The reason is that compliance arising naturally in a transmission may not be very linear; often it stiffens with increasing force, and the stiffness may vary with position. The intentionally-added spring can be designed to be linear over the force range of interest and softer than the drive compliance, so that its behavior is dominant and the relative modeling error caused by the transmission nonlinearity is much reduced.

If a more compliant drive is used with any of the controllers discussed in Chapter 3, the positioning accuracy of the end-effector decreases. Since significant load-dependent errors already exist when these controllers are used with relatively stiff drives, increased error could be a serious drawback. This is addressed in Section 4.2 by the addition of joint-position sensing, which is a necessary component of the new controller design. Joint sensing allows joint position errors to be rejected, but since it involves adding a non-colocated sensor, a more sophisticated control approach is required.

When joint position sensing is employed to reduce endpoint position errors, there arises a distinction between drive compliance and the compliance outboard of the joint

position sensor. Such outboard compliance may be due to link bending, or flexible elements in the force sensor or end-effector. Increasing outboard compliance provides the same benefits for smoothness and contact behavior as increasing drive compliance, but positioning accuracy would be lost unless some kind of absolute endpoint position sensing is used (as suggested by Cannon et al.^[9,10]). This possibility is not pursued here. Compliance outboard of the joint simply contributes to the effective environmental compliance.

4.1.4 Actuator Force Requirements

Mechanical drive decoupling has one effect which can be a serious practical drawback. Operation of a manipulator above its drive frequency ω_D , as implied by drive decoupling, incurs a penalty in power consumption and force requirements. At high frequencies a given endpoint trajectory requires higher-amplitude actuator motion when drive compliance is increased. More power is consumed through acceleration of the actuator inertia and dissipation from actuator-axis friction. Considering the inertial effect alone, the peak actuator force required to track a sinusoidal endpoint position trajectory of frequency ω with a compliant-drive manipulator is:

$$U_{\max,\text{compliant}} = \left[1 - \frac{m_L + m_P}{m_A + m_L + m_P} \left(\frac{\omega^2}{\omega_D^2} \right) \right] U_{\max,\text{rigid}} \quad (4.1.1)$$

where $U_{\max,\text{rigid}}$ is the corresponding peak force for a rigid-drive manipulator.

For high-bandwidth operation, $\omega > \omega_D$ and $U_{\max,\text{compliant}}$ becomes larger than (and 180° out of phase with) $U_{\max,\text{rigid}}$. Lowering the drive frequency ω_D increases peak force requirements. This is compounded by the fact that the controller improvements discussed in the next section allow a higher overall bandwidth of stable operation. Taking advantage of this also increases force requirements. Actuator force constraints may limit the amount of decoupling compliance it is desirable to introduce for a given intended bandwidth of operation.

The controllers described in Chapter 3 are restricted to a usable bandwidth below the drive natural frequency, so their peak force requirements are essentially those of a rigid manipulator. If drive stiffness is lowered by a factor of 10 and operating bandwidth increased to twice the original drive frequency, eqn. (4.1.1) predicts that peak forces required would increase by a factor of 19 (for tracking of a sinusoidal trajectory at the bandwidth limit). The drive decoupling alone, with no increase in bandwidth, gives a peak force increase of a factor of 4.

Equation (4.1.1) gives only a rough indication of this effect; it is difficult to predict on analytical grounds the level of actuator force required to follow an arbitrary reference trajectory, while rejecting disturbances and reacting to environmental contact. Numerical simulation is a better tool for studying actuator force levels. The analysis in this chapter will concentrate on the performance benefits obtainable when full advantage is taken of the stability and bandwidth limits of the new controller design, with a decoupled drive. Simulations in the next chapter explore several ways to reduce actuator force requirements.

4.2 Joint/Actuator Control Loops

In all of the controllers described in Chapter 3, bandwidth was limited to below the drive resonance frequency ω_D , where joint and actuator positions are essentially the same except for a DC offset. These control laws implicitly assume that forces applied by the actuator are instantly conveyed to the joint. The phase shift due to drive compliance is “unexpected” and causes instability.

The controller proposed here is based on a compliant-drive model and allows operation above the drive resonance frequency ω_D .

4.2.1 Controller Structure

This approach uses a position sensor on the joint in addition to one on the actuator. The desired joint trajectory is derived as before from the input reference trajectory x_{ref} , modified by the measured interaction force F and commanded stiffness $(d_T s + k_T)$. But now, a joint control loop uses feedback from the joint sensor to compute the force command required to correct the joint position error. A model of drive compliance is then used to calculate the actuator position required to provide the desired joint force. The resulting actuator trajectory is then tracked by the inner control loop. This is illustrated by the following block diagram:

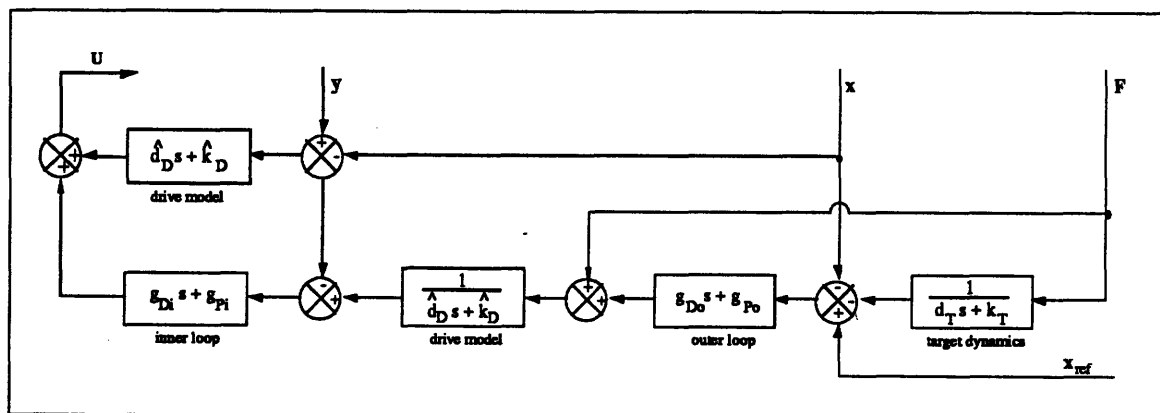


Figure 4.1: Joint/Actuator Controller Block Diagram

The control law is:

$$u = (\hat{d}_D s + \hat{k}_D)(y - x) + \quad (4.2.1)$$

$$(g_{Da} s + g_{Pa}) \left\{ \left(\frac{1}{\hat{d}_D s + \hat{k}_D} \right) \left[F + (g_{Dj} s + g_{Pj}) \left(x_{\text{ref}} - \frac{F}{d_T s + k_T} - x \right) \right] + x - y \right\}$$

The first term is feedforward of the estimated drive compression force. The actuator-loop gains are g_{Da} and g_{Pa} , and the joint-loop gains are g_{Dj} and g_{Pj} . The model drive parameters are \hat{d}_D and \hat{k}_D , which are estimates of the real parameters d_D and k_D .

This controller differs from the compliance-modeling approaches discussed in Section 2.4.1 in several ways. One difference is the use of a joint position sensor. Another is the “classical control” approach of nested control loops, rather than a “modern control” state-variable formulation. For this approach to be valid, the interaction between the control loops must be small. This requires the bandwidth of the inner (actuator) loop to be higher than that of the outer (joint) loop, and the joint and actuator variables to be mechanically decoupled. In this case, the latter condition means that inner-loop bandwidth must be much higher than the drive natural frequency. For disturbance rejection, controllers already in use (e.g. the reference design) have much higher inner-loop bandwidths than overall bandwidths, so the first assumption is easily satisfied. As described in the previous section, decoupling of the joint and actuator positions typically requires the addition of extra compliance to the transmission.

The advantage of this approach is a fairly simple controller structure, with no high-order estimation or differentiation required. No unstable poles or zeroes are cancelled. The gains have straightforward physical interpretations and the loops can be tuned independently. A simple criterion for stability can be derived to give the designer a logical way of picking gains.

This controller requires additional drive compliance for another reason as well. The position variables x and y are subtracted in the control law to determine the force being transmitted by the drive. Since the position sensors have finite resolution, subtracting two nearly equal numbers would provide very poor open-loop joint-force resolution. Compliance increases the magnitude of the position difference for a given force, refining this resolution.

Separate actuator/joint control loops are intended to improve performance in the areas of stability and bandwidth. Use of a model of drive compliance allows stable

operation above ω_D and permits the closed-loop stiffness to closely approximate the target value, even when this exceeds drive stiffness.

To see how force-control stability range can improve, examine the problem from the point of view of the joint, which in our model is rigidly connected to the force sensor. If an ideal actuator is also coupled rigidly to the joint, no instability occurs, even with a stiff environment and soft commanded compliance. However, if a transmission (or actuator) with limited bandwidth is used, the phase shift introduced into joint force commands as frequencies approach the bandwidth limit causes instability.

With the Joint/Actuator controller design, transmission dynamics are cancelled out by the internal model used by the inner loop. There is still, however, a limiting bandwidth for joint force commands--the inner-loop bandwidth itself. As frequencies approach this bandwidth limit, phase shift is introduced and instability again results.

As far as the joint loop is concerned, drive dynamics have been replaced by the inner-loop dynamics. Since the inner loop typically has a higher bandwidth and much higher damping ratio than the drive compliance, instability is postponed. The effect is similar to stiffening the drive, but is accomplished in software rather than hardware.

A consequence of this, which will be explored later, is that the effect of inner-loop gain on stability is reversed for the Joint/Actuator controller: with the controllers of Chapter 3, an increase of inner-loop bandwidth is destabilizing; with the new controller, stability is improved as inner-loop bandwidth increases.

4.2.2 Comparison with Force-Controlling Inner Loops

The Joint/Actuator control scheme bears some similarity to the direct feedback force control scheme suggested by Luh, Fisher, and Paul^[32]. In their system, a force sensor is used in the transmission to measure the force actually applied to the joint. A high-gain force feedback loop is then used to achieve the desired force and reject actuator-axis force disturbances. No model of drive compliance is required and no mechanical drive decoupling is needed.

By contrast, the Joint/Actuator controller described here forms its estimate of the joint force error from measurements of relative joint and actuator position and a model of drive compliance. The high-gain inner loop then acts to reduce this position error. There are some fundamental advantages to this approach.

First, the sensor on which the inner loop is based is actuator position, which is colocated. Position sensors such as optical encoders introduce no significant low-

frequency dynamics of their own; this permits high gains. Force sensors usually include a flexible element which can be destabilizing at high gains.

Second is the fact that most of the disturbance forces acting on the actuator are position- or velocity-related: friction, stiction, and cogging. If a model is available for these disturbances, a position-based inner loop can use feedforward compensation to reduce their effect *before* they cause position errors. As noted by Canudas et al.^[11] the parameters of the friction disturbances change somewhat with operating conditions; they propose an adaptive scheme for feedforward compensation. Even without this level of effort, significant improvements can be obtained by feeding forward easily-calculated forces which underestimate the full value of the disturbance (see for instance [23]). A force-based inner loop cannot predict these disturbances and must rely on feedback to correct errors after they occur.

Third is the issue of force sensor placement. The sensor can be placed near the actuator axis or near the joint axis, the difference being whether transmission dynamics are enclosed in the inner loop. If placed on the actuator axis, disturbance forces arising in the transmission cannot be rejected by the inner loop, since the input force to the transmission is being controlled instead of the output force. If placed on the joint axis, inertia and compliance in the transmission makes the sensor non-colocated, limiting gains and eliminating the principal advantage of an inner loop. Luh, Fisher, and Paul used actuator-axis sensor placement in their experiments.

The following simple model shows that the position-based inner loop of the Joint/Actuator controller, combined with drive decoupling, can reject disturbance forces arising in the transmission, while a force-sensing loop with a sensor on the actuator axis cannot. Consider a transmission model consisting of a massless node upon which friction (or some other disturbance) acts. This is connected to the actuator mass m_A by a spring representing transmission compliance k_A and connected to the link mass m_L by the spring k_J added to decouple the drive, as shown in Figure 4.2:

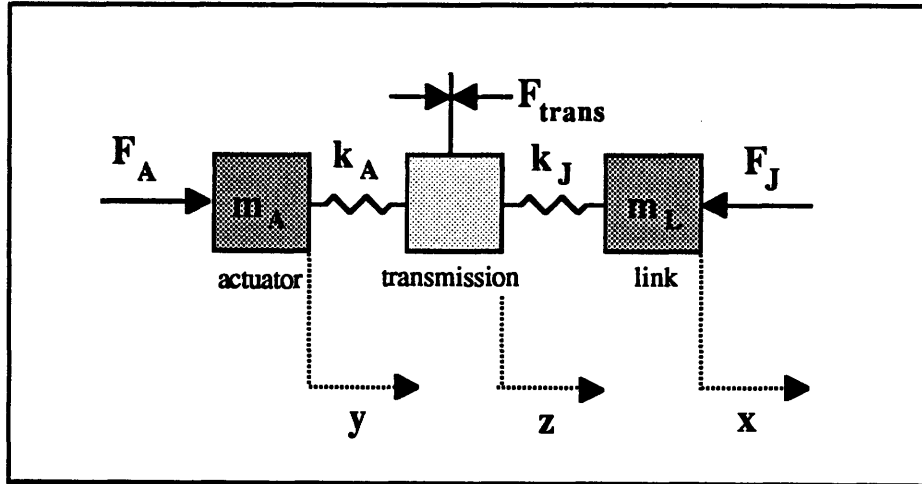


Figure 4.2: Transmission Disturbance Model

The compression force in the transmission compliance is the actuator force F_A and the force in the decoupling spring is the joint force F_J , given by:

$$F_A = k_A(y - z) \quad (4.2.2)$$

$$F_J = k_J(z - x)$$

In equilibrium the forces must sum to zero at the transmission node, so:

$$F_A - F_J = F_{trans} \quad (4.2.3)$$

Assume the commanded joint force is zero. If the magnitude of the transmission force disturbance is F_{trans} then the actuator can be producing zero force while the joint force equals the full disturbance. From eqn. (4.2.3):

$$F_A = 0 \Rightarrow F_J = F_{trans} \quad (4.2.4)$$

The force-sensing inner loop with actuator-axis sensor placement thus cannot reject *any* transmission force disturbance. The Joint/Actuator controller, if commanded to produce zero joint force, drives the relative actuator and joint positions to $x - y = 0$. Equations (4.2.2) and (4.2.3) then imply:

$$F_J = F_{\text{trans}} \left(\frac{k_J}{k_J + k_A} \right) \quad (4.2.5)$$

Thus, if the drive-decoupling spring k_J is N times softer than the initial drive stiffness k_A , the Joint/Actuator inner-loop attenuates transmission disturbance forces by a factor of $(N + 1)$. The Joint/Actuator controller therefore rejects force disturbances in the transmission, while the force-sensing loop cannot.

Another implication of the above result is that the decoupling spring should be inserted between the transmission and joint, rather than between transmission and actuator. In other words, the friction-introducing part of the transmission should be placed on the actuator rather than the joint. This minimizes the force error caused by transmission disturbances.

Joint-axis force sensor placement is the other alternative for a force-controlling inner loop, and this allows rejection of transmission disturbances. However, any dynamics between the actuator and control sensor can be destabilizing to the force-control loop. In the Joint/Actuator control scheme such dynamics cause error but not instability. The following example clarifies this.

Figure 4.3a shows a simplified version of the drive model, with the link fixed and drive damping neglected. A real transmission, of course, is not a purely massless spring. In gearboxes, for example, the compliance is due to shaft and tooth bending, and the gears themselves have inertia. A more accurate model would include an inertia for each gear, and springs connecting them in series. A step in this direction is shown in Figure 4.3b, with one intermediate inertia.

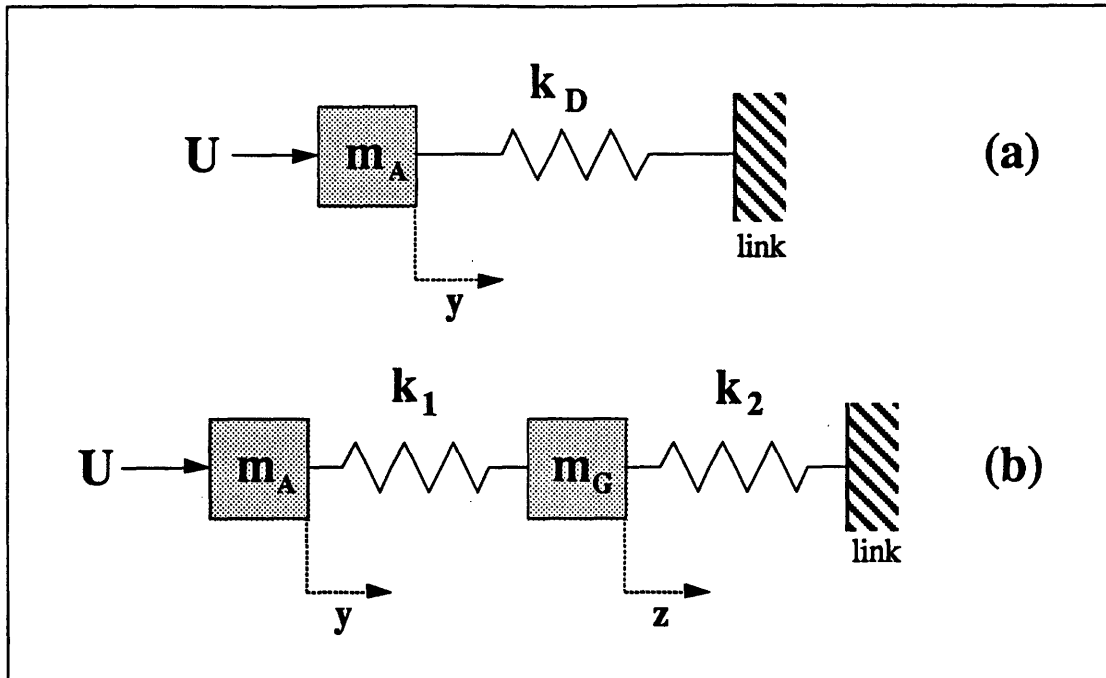


Figure 4.3: Transmission Models

To preserve the same overall drive compliance as in Figure 4.3a, the following equation must hold:

$$\frac{1}{k_D} = \frac{1}{k_1} + \frac{1}{k_2} \quad (4.2.6)$$

The dynamics of the intermediate mass m_G are given by:

$$m_G \ddot{z} = k_1 (y - z) - k_2 z \quad (4.2.7)$$

If a force sensor is installed at the joint-axis end of the transmission, it will measure the compression of the second spring:

$$F = k_2 z \quad (4.2.8)$$

The transfer function from actuator position y to measured force F is then:

$$\frac{F(s)}{y(s)} = \frac{k_1 k_2}{m_G s^2 + k_1 + k_2} \quad (4.2.9)$$

Thus there are now second-order dynamics between the actuator and control sensor, with natural frequency:

$$\omega_n = \sqrt{\frac{k_1 + k_2}{m_G}} \quad (4.2.10)$$

This is a limit on the maximum bandwidth of a simple inner control loop using joint-axis force sensing. To estimate this numerically, use the value of $k_D = 1000$ N/m from Table 3.1, and assume that $m_G = .1 m_A$ and that $k_1 = k_2$. This then gives $\omega_n = 32$ Hz, a considerable restriction on inner-loop bandwidth; the reference design controller of Chapter 3 used 50 Hz.

The simple example above gives a first-order illustration of the effects of distributed inertia and compliance in the transmission. In the limit of a continuum model, a force change would propagate as a wave from the actuator to the joint. It is the delay due to finite wave speed that causes instability with high gains. An inner loop based directly on actuator position is colocated and does not suffer from this problem; it can achieve higher bandwidth.

One other somewhat similar force control system in the literature is the Whole Arm Manipulation (WAM) concept of Salisbury^[51]. Force sensing in the transmission is used there as well. In this case, however, transmission force sensors take the place of the wrist-mounted force sensor instead of replacing the inner-loop position sensors. A velocity-based inner loop is still used around each actuator. The reasons for this derive from the WAM idea, which is to be able to use the entire arm surface for manipulation, rather than just the end-effector. The merits of this philosophy will not be pursued here; the manipulators studied in this thesis are of the conventional type simply to make the results relevant to a wider body of research. It should be noted, however, that for conventional tasks, placing the force sensor inboard of the link mass is detrimental to contact behavior. This is explained more fully in Section 4.3.4.

4.3 Performance Evaluation

In this section the ideas presented above are explored and evaluated according to the performance criteria adopted in Section 3.1.2. To facilitate comparison with the controller analysis in Chapter 3, the parameter values of the following table are adopted as a standard example:

Parameter	Reference Design Value	Joint/Actuator Design Value
m_A	1 kg	1 kg
m_L	1 kg	1 kg
m_P	.01 kg	.01 kg
d_D	5 Ns/m	.5 Ns/m
d_E	2 Ns/m	2 Ns/m
d_T	60 Ns/m	60 Ns/m
k_D	1000 N/m	100 N/m
k_E	20000 N/m	20000 N/m
k_T	600 N/m	600 N/m
drive frequency	$\omega_D = 5.03$ Hz	$\omega_D = 1.59$ Hz
environment frequency	$\omega_E = 22.5$ Hz	$\omega_E = 22.5$ Hz
payload frequency	$\omega_P = 225$ Hz	$\omega_P = 225$ Hz
inner-loop b/w	$\omega_i = 50$ Hz	$\omega_a = 50$ Hz
inner-loop ζ	$\zeta_i = .707$	$\zeta_a = .707$
outer-loop b/w	$\omega_o = 6$ Hz	$\omega_j = 10$ Hz
outer-loop ζ	$\zeta_o = .707$	$\zeta_j = .707$

Table 4.1: Controller Comparison Parameters

The values for the reference design match those used in Chapter 3. Differences for the Joint/Actuator design (highlighted by bold-face type), are the drive parameters (reduced to provide mechanical decoupling), and the outer-loop bandwidth. The latter is limited by the requirement of stability, and one of the chief benefits of the new design is its ability to achieve a higher bandwidth while retaining stability. This is discussed in the next section.

4.3.1 Stability Range

This section characterizes the stability limits of the Joint/Actuator design approach, and contrasts them with those of the reference design. The effects of drive decoupling on the FFTF are first examined, then the effects of the new controller structure on the RFTF.

Forward Force Transfer Function

Increasing drive compliance changes the FFTF. For the Joint/Actuator controller example the drive is decoupled by reducing k_D and d_D by a factor of ten. Figures 4.4 and 4.5 show the new FFTF. The rigid-drive FFTF is also plotted for reference.

The major differences between these plots and those in Figures 3.5 and 3.6 are the lowering of the actuator-drive resonance frequency ω_D from 5 Hz to 1.6 Hz, the increase in magnitude of that peak from 17.2 db to 26.1 db, and the reduction of the magnitude of the link-environment resonance peak at ω_E from -0.3 db to -9.2 db. The phase margin at ω_D falls from 25° to 8° and the gain margin at ω_E increases from 0.3 to 9.2 db.

As seen in Chapter 3, instability with the reference design controller occurs when a sharp dip in the RFTF response near ω_E uses up the gain margin there, and the RFTF does not contribute sufficient phase lead at that frequency to bring total phase above -180°.




Figure 4.4: Decoupled-Drive FFTF Magnitude




Figure 4.5: Decoupled-Drive FFTF Phase

Return Force Transfer Function

From the control law eqn. (4.2.1) and eqn. (3.3.1) are derived the following transfer functions:

$$A(s) = \left(\frac{g_{Da}s + g_{Pa}}{\hat{d}_D s + \hat{k}_D} \right) \left(1 - \frac{g_{Dj}s + g_{Pj}}{d_{Ts} + k_T} \right)$$

$$B(s) = (g_{Da}s + g_{Pa}) - (\hat{d}_D s + \hat{k}_D) - \left(\frac{g_{Da}s + g_{Pa}}{\hat{d}_D s + \hat{k}_D} \right) (g_{Dj}s + g_{Pj}) \quad (4.3.1)$$

$$C(s) = (\hat{d}_D s + \hat{k}_D) - (g_{Da}s + g_{Pa})$$

$$R(s) = \left(\frac{g_{Da}s + g_{Pa}}{\hat{d}_D s + \hat{k}_D} \right) (g_{Dj}s + g_{Pj})$$

Eqn. (3.4.1) then gives the RFTF for this controller:

$$\begin{aligned} H(s) = & \left(\frac{g_{Da}s + g_{Pa}}{\hat{d}_D s + \hat{k}_D} \right) \left\{ \left(\frac{\hat{d}_D s + \hat{k}_D}{d_D s + k_D} \right) \left[1 + \frac{m_L s^2}{m_P s^2 + d_E s + k_E} \right] - 1 \right. \\ & + \left. \left(\frac{g_{Dj}s + g_{Pj}}{d_{Ts} + k_T} \right) \left[1 + \frac{d_{Ts} + k_T}{m_P s^2 + d_E s + k_E} \right] \right\} \quad (4.3.2) \\ & - \left(\frac{\hat{d}_D s + \hat{k}_D}{d_D s + k_D} \right) \left[1 + \frac{m_L s^2}{m_P s^2 + d_E s + k_E} \right] \end{aligned}$$

With the assumption of accurate modeling, this becomes:

$$\begin{aligned} H(s) = & \left(\frac{g_{Da}s + g_{Pa}}{d_D s + k_D} \right) \left[\frac{m_L s^2 + g_{Dj}s + g_{Pj}}{m_P s^2 + d_E s + k_E} + \frac{g_{Dj}s + g_{Pj}}{d_{Ts} + k_T} \right] \\ & - \left(\frac{m_L s^2}{m_P s^2 + d_E s + k_E} \right) - 1 \quad (4.3.3) \end{aligned}$$

Figures 4.6 and 4.7 show the Bode plot for this RFTF, using the parameter values from Table 4.1.

The shape of the RFTF plots are basically similar to those of the other stiffness controllers with inner loops. Comparison of equations (3.4.11) and (4.3.3) shows that the RFTF's have the same form. As long as actuator bandwidth is reasonably high, the feedforward terms in each equation can be neglected for stability analysis.

As before, the effect of the feedback term is to give approximately 180° phase lead between a frequency ω_R and the payload-environment resonance frequency ω_P . The frequency ω_R of the first resonance is different in this case, however. Using the same reasoning as in Section 3.4.4, ω_R is found to be:

$$\omega_R = \sqrt{\frac{g_D k_E}{m_L d_T} + \omega_j^2} \quad (4.3.4)$$

where ω_j is the bandwidth of the joint loop. For the parameters of Table 4.1, this gives $\omega_R = 29.2$ Hz.

Note that ω_R for this controller is independent of drive parameters k_D and d_D . Another important difference is that for this controller the depth of the resonance magnitude dip at ω_R is determined by the (user-selected) damping of the joint loop instead of the (often very low) damping of the drive and environment. As a result the dip is much broader and shallower. In the reference design controller, the depth of this dip sets the stability limit by causing the open-loop transfer function gain to cross over unity there. In the new controller, the gain crossover frequency (and stability limit) is not set by the dip at ω_R but by the general rolloff of the FTF gain at higher frequencies.

This can be seen in the open-loop transfer function plotted in Figures 4.8 and 4.9. The gain crossover frequency in this case is 74 Hz. The phase margin is 51.1°, compared with 16.4° in the reference design case, indicating significantly more robust stability.




Figure 4.6: Joint/Actuator Controller; RFTF Magnitude




Figure 4.7: Joint/Actuator Controller; RFTF Phase




Figure 4.8: Joint/Actuator Controller; Open-Loop Magnitude




Figure 4.9: Joint/Actuator Controller; Open-Loop Phase

At frequencies above ω_R the following approximations hold, from equations (4.3.3) and (3.3.5):

$$H(s) \approx \left(\frac{g_D a s}{d_D s + k_D} \right) \left[\frac{m_L s^2}{m_P s^2 + d_E s + k_E} \right]$$

$$G(s) \approx \frac{(m_P s^2 + d_E s + k_E)(d_D s + k_D)}{m_A m_L s^4} \quad (4.3.5)$$

$$\therefore G(s)H(s) \approx \frac{g_D a}{m_A s}$$

The gain crossover frequency ω_{GC} is thus approximately given by:

$$\omega_{GC} \approx \frac{g_D a}{m_A} = 2\zeta_a \omega_a \quad (4.3.6)$$

This approximation gives a value of 71 Hz in this case; close to the actual value of 74 Hz.

Since the phase of the open-loop transfer function rises above -180° near the frequency ω_R , an approximate criterion for stability can be stated $\omega_{GC} > \omega_R$, which becomes, using equations (4.3.4) and (4.3.5):

$$4\zeta_a^2 \omega_a^2 > \omega_j^2 + \frac{2\zeta_j \omega_j k_E}{d_T} \quad (4.3.7)$$

This can be rearranged with controller gain parameters on one side, showing the limitation on commanded behavior due to instability:

Joint/Actuator Stability Criterion:

$$\frac{d_T}{k_E} > \frac{2\zeta_j \omega_j}{4\zeta_a^2 \omega_a^2 - \omega_j^2} \quad (4.3.8)$$


Two assumptions were used in deriving this stability criterion:

- i) $d_T \omega_R > k_T$, which is true for any reasonably-damped desired behavior, and
- ii) $\omega_P > \omega_{GC}$, which means the environment is primarily springlike, rather than inertial

The first assumption allowed approximation of the term $(d_T s + k_T)$ by $d_T s$ at $s = i\omega_R$, and the second allows us to approximate $(m_{PS} s^2 + d_E s + k_E)$ by k_E at $s = i\omega_{GC}$, where we also assume environmental damping d_E is negligible.

Since it is an approximation, the stability criterion must be applied conservatively, but should give a good indication of the effects of varying each parameter. From this criterion it is evident that actuator bandwidth ω_a should exceed joint bandwidth ω_j . It can also be seen that increasing the commanded stiffness or softening the environment improves stability, as usual.

To test assumption i), the tradeoff between lowest attainable values of d_T and k_T is plotted in the following figure (for parameter values from Table 4.1), using the Routh criterion on the closed-loop characteristic polynomial to determine stability:



The left-hand stability boundary (solid line) of the shaded area is nearly vertical; there is little dependence on k_T , so assumption i) above is justified here. Comparison with Figure 3.30 reveals improved stability range for the new controller, using the same actuator-loop gains. The lowest pure damping value attainable by the reference design was 54 Ns/m, versus 38 Ns/m for the new design. Note also that $\omega_j = 10$ Hz for the new controller and $\omega_0 = 6$ Hz for the reference design in this comparison. Reducing ω_j to 6 Hz

would give an even larger stability range improvement, allowing pure damping values as low as 26 Ns/m.

Large payload mass would invalidate assumption ii) above. In this case the proper approximation becomes $(m_P s^2 + d_E s + k_E) \approx m_P s^2$ at $s = i\omega_{GC}$. The gain crossover frequency is then given by:

$$\omega_{GC} \approx 2\zeta_a \omega_a \left(\frac{m_L}{m_L + m_P} \right) \left[1 + \frac{\zeta_j \omega_j}{\zeta_T \omega_T} \right] \quad (4.3.9)$$

where ω_T is defined as the natural frequency of the payload mass with the commanded stiffness k_T , and ζ_T is the corresponding damping ratio:

$$\omega_T \equiv \sqrt{\frac{k_T}{m_P}} \quad (4.3.10)$$

$$\zeta_T \equiv \frac{d_T}{2 \sqrt{m_P k_T}}$$

In the large-payload case, the phase of the open-loop transfer function rises above -180° at ω_R or ω_P , whichever is lower. The stability criterion becomes:

$$2\zeta_a \omega_a \left(\frac{m_L}{m_L + m_P} \right) \left[1 + \frac{\zeta_j \omega_j}{\zeta_T \omega_T} \right] > \text{Min} \left(\sqrt{\omega_j^2 + \frac{2\zeta_j \omega_j k_E}{d_T}}, \omega_P \right) \quad (4.3.11)$$

Assuming $\zeta_T \approx \zeta_j$, it can be shown that sufficient conditions for this stability criterion to be satisfied are:

- the no-payload stability criterion (4.3.8) above, and
- $g_{Pj} > \text{Max} (k_T, \sqrt{k_E k_T})$

Since high gains and low commanded stiffness are likely to be desirable operating conditions, the previous stability criterion (4.3.8) will imply stability in this case as well. In other words, if a manipulator in contact with the environment is stable without a payload, it will usually be stable with one. As is the case for other controller designs, stability with a stiff springlike environment is more difficult to obtain than with an inertial environment.

In summary, the Joint/Actuator design shows improved stability tradeoffs over the reference design, allowing a wider range of commanded behavior. A simple approximate

criterion for stability was derived for the new controller (eqn. 4.3.8), which allows the choice of controller parameters to suit a particular task.

4.3.2 Control Bandwidth

The closed-loop tracking quality and bandwidth can be evaluated using the closed-loop admittance given by equation (3.4.8). With the transfer functions for the new controller (eqns. 4.3.1) this is:

$$\frac{x(s)}{F(s)} = \frac{m_A s^2 (d_T s + k_T) + (g_{D_a} s + g_{P_a})(g_{D_j} s + g_{P_j})}{(d_T s + k_T) [m_A s^2 (m_L s^2 + d_D s + k_D) + (g_{D_a} s + g_{P_a})(m_L s^2 + g_{D_j} s + g_{P_j})]} \quad (4.3.12)$$

In the limit of very high actuator-loop gains g_{D_a} and g_{P_a} this becomes:

$$\frac{x(s)}{F(s)} \approx \left(\frac{g_{D_j} s + g_{P_j}}{m_L s^2 + g_{D_j} s + g_{P_j}} \right) \left(\frac{1}{d_T s + k_T} \right) \quad (4.3.13)$$

It can be seen by comparison with eqn. (3.4.15) that the joint loop in the new controller is directly analogous to the filter in the reference design. When the loop is closed, the effect is to filter the desired admittance $\left(\frac{1}{d_T s + k_T} \right)$ by a factor $\left(\frac{g_{D_j} s + g_{P_j}}{m_L s^2 + g_{D_j} s + g_{P_j}} \right)$ of bandwidth ω_j , whereas the reference design directly filters the measured forces by the factor $\left(\frac{g_{D_0} s + g_{P_0}}{m_L s^2 + g_{D_0} s + g_{P_0}} \right)$. In determining bandwidth, therefore, joint-loop gains play a similar role to the filter gains of the reference design.

Equation (4.3.13) also shows that with a high-gain inner loop, the overall bandwidth of the new controller is simply that of the joint loop. Since drive dynamics are absent from this equation, the achieved admittance should be accurate throughout the bandwidth.

Figure 4.11 plots the achieved vs. target admittance for the Joint/Actuator controller, using parameters from Table 4.1.




Figure 4.11: Joint/Actuator Controller; Closed-Loop Admittance

Compared with Figure 3.25 for the reference design, this figure shows much better tracking of the target admittance. No steady-state error exists and the achieved admittance does not deviate markedly from the target admittance until the frequency exceeds 20 Hz. This was achieved using the same inner-loop gains as the reference design, and with a greater margin of stability.

4.3.3 Smoothness

This section compares the new controller to the reference design using some criteria for trajectory smoothness, primarily the degree of rejection of actuator force and position disturbances. Then, since dry friction is poorly modeled as a linear disturbance, numerical simulation is used to evaluate the effects of drive compliance and controller bandwidth on smoothness when Coulomb friction and stiction are present.

Force Disturbance Rejection

The force disturbance transfer function is given by equation (3.3.4). For the Joint/Actuator controller it is:

$$\left[\frac{F(s)}{D(s)} \right]_{J/A} = \frac{(d_T s + k_T)(m_P s^2 + d_E s + k_E)(d_D s + k_D)}{m_A s^2 (d_T s + k_T) [(m_P + m_L) s^2 + (d_E + d_D) s + k_E + k_D]} \dots \quad (4.3.14)$$

$$\dots + (g_{D_a} s + g_{P_a}) [m_L s^2 (d_T s + k_T) + (g_{D_j} s + g_{P_j})(m_P s^2 + (d_E + d_T) s + k_E + k_T)]$$

Low-frequency behavior is particularly important. In the limit it becomes:

$$\left[\frac{F(s)}{D(s)} \right]_{J/A} \rightarrow \frac{k_D k_E k_T}{g_{P_a} g_{P_j} (k_E + k_T)} \quad \text{as } s \rightarrow 0 \quad (4.3.15)$$

For the reference design controller this limit becomes (from eqn. 3.4.18):

$$\left[\frac{F(s)}{D(s)} \right]_{REF} \rightarrow \frac{k_D k_E k_T}{g_{P_i} (k_D k_E + k_D k_T + k_E k_T)} \quad \text{as } s \rightarrow 0 \quad (4.3.16)$$

When $k_E > k_D > k_T$, a typical desirable operating condition, the ratio of these transfer functions becomes:

$$\frac{(F/D)_{REF}}{(F/D)_{J/A}} \approx \frac{g_{P_j}}{k_D} \quad (4.3.17)$$

assuming the inner-loop gains are the same in each case. The k_D appearing in this equation is the drive stiffness of the Joint/Actuator controller rather than the reference design, if the two differ. In this approximation, therefore, disturbance rejection with the new controller is improved over the reference by increasing joint-loop gains or by adding drive compliance

(lowering k_D). A net improvement occurs if the joint loop has higher bandwidth than the drive resonance frequency (which is a primary objective of the new controller anyway).

The following figure compares the magnitude of the actuator-axis force disturbance transfer function for the Joint/Actuator controller to that of the reference design, using the parameter values from Table 4.1:




Figure 4.12: Force Disturbance Rejection Comparison

The comparison reveals an improvement over the reference design of approximately 30 db (a factor of 32) in disturbance rejection throughout the controller's bandwidth. The improvement at higher frequencies is approximately 20 db (a factor of 10).

Position Disturbance Rejection

Limit cycles due to backlash or other causes can behave more like actuator position disturbances than force disturbances, in the sense that turning up inner-loop gains may not reduce their amplitude. Also, in practice there are limits to how high gains can be increased, and limits to sensor resolution, which combine to give the minimum-attainable actuator position error. This section investigates the effect of this limit on force control accuracy.

The transfer function from actuator position y to measured force F is given by:

$$\frac{F(s)}{y(s)} = \frac{(m_P s^2 + d_E s + k_E)(d_D s + k_D)}{(m_P + m_L)s^2 + (d_E + d_D)s + k_E + k_D} \quad (4.3.18)$$

Since no control input occurs in the system between the actuator and the manipulator tip, this transfer function depends only on the mechanical parameters; controller design is irrelevant. With a rigid-drive manipulator, the transfer function approaches $(m_P s^2 + d_E s + k_E)$. If the environment is stiff, large forces result from small position displacements. The manipulator designer can reduce this effect only by softening the effective environment (with a compliant covering on the manipulator, for instance), or by reducing the drive stiffness.

The equation indicates, however, that in the usual case of an environment more rigid than the drive, it is far more effective to soften the drive than to soften the environment. Figure 4.13 illustrates the advantage obtained by reducing drive stiffness by a factor of 10, from the value used for the reference design example to that of the Joint/Actuator (with drive decoupling) design example. Also shown is the effect of reducing the environment stiffness by a factor of 10.

Note that, as in all of the “disturbance rejection” plots, the curves show the magnitude of the transfer function from disturbance to output; therefore lower values are better. In the present case of a transfer function from position disturbance to output force, the actual values depend on the units used: here 0 db indicates a 1 N response to a 1 m disturbance. As usual, relative comparisons of the values on different curves (expressed in db) represent dimensionless ratios.




Figure 4.13: Position Disturbance Rejection Comparison

Softening the drive improved rejection by 20 db at nearly all frequencies, while softening the environment gave a much smaller improvement (about 4 db) over most of the manipulator's bandwidth. Installing a compliant covering on a manipulator reduces its versatility, by increasing its bulk and limiting the stiffness it can attain. Since this is not a very effective means of improving actuator position disturbance rejection, the alternative of adding drive compliance is more attractive.

This example shows that adding drive compliance can significantly reduce the effect of actuator position disturbances on force control, contributing to smoothness of response.

Force Dynamic Range

The maximum static force available is the same for both designs, though in operation the force reserved for dynamic disturbance rejection will depend on factors such as bandwidth. The minimum force increment can be no smaller than:

- i) the force corresponding to the least significant bit of control signal resolution
- ii) the magnitude of the disturbance forces times the steady-state magnitude of the force disturbance transfer function, or
- iii) the magnitude of the position disturbances times the steady-state magnitude of the position disturbance transfer function, or in the case of limit cycles, the magnitude of the transfer function at the limit cycle frequency.

In the example of Table 4.1, the above transfer functions indicate the Joint/Actuator design would have between 10 and 30 times the static force dynamic range of the reference design, assuming that limitation i) is avoided by proper selection of electronic components.

Stick-Slip Simulation

Linear transfer function analysis is valuable but doesn't tell the whole story when a disturbance is highly nonlinear, as in the case of dry friction. Simulation provides a means of investigating the effects of nonlinear disturbances. It does not replace analysis because it is difficult to generalize quantitatively, but it can give useful qualitative information and compare performance in specific cases.

Friction is known to cause "stick-slip" behavior in mechanisms. This is antithetical to the goal of smoothness of response. Since a major objective of the Joint/Actuator controller design is to improve smoothness, a simplified problem is simulated here to investigate the controller's effectiveness in suppressing this specific type of disturbance.

This simulation compares performance of the position-controlling parts of the controllers only, so the results do not depend on environmental parameters. The desired trajectory is assumed to be a simple ramp in free space, starting from rest:

$$\begin{aligned}x_T(t) &= Vt, t > 0 \\x_T(0) &= \dot{x}_T(0) = 0\end{aligned}\tag{4.3.19}$$

The target velocity value used is 5 cm/sec. In each case the controller is assumed to comprise an outer (joint) loop and an inner (actuator) loop, with specified bandwidths and damping ratios ($= .707$ in each case). These and the drive stiffness are selected to correspond to the reference design example and the Joint/Actuator example used above. An intermediate case shows the effect of mechanical decoupling alone. The following table gives the values used:

Parameter	Reference Design Value	Mechanically Decoupled Value	Joint/Actuator Design Value
k_D	1000 N/m	100 N/m	100 N/m
ω_{inner}	50 Hz	50 Hz	50 Hz
ω_{outer}	6 Hz	6 Hz	10 Hz

Table 4.2: Stick-Slip Simulation Parameters

The disturbance forces are Coulomb friction and stiction. Very high values of these forces were needed ($F_{\text{Coulomb}} = 100 \text{ N}$, $\Delta F_{\text{stiction}} = 500 \text{ N}$) to produce observable errors with the control gains assumed; as seen above, both example controllers have in excess of 45 db force disturbance rejection. Although this level of friction is unlikely, the same disturbance values are used for each case so the relative comparison is still valid. Many practical manipulators use lower-bandwidth inner loops than in these examples, and can exhibit this behavior with prevailing levels of friction. A factor of ten lower in bandwidth brings the required friction levels to values typically seen.

The simulation was performed using a 4th-order Runge-Kutta scheme with a time step of .001 second.

Figure 4.14 shows the joint-position response in the reference design case. Figure 4.15 shows the actuator-position trajectory. Stick-slip behavior dominates the actuator plot, and its effects are evident in the tracking error of the joint plot.




Figure 4.14: Reference Design; Joint-Position Response




Figure 4.15: Reference Design; Actuator-Position Response

Increasing the drive compliance, with no other change to the system, produces the results shown in Figures 4.16 and 4.17. Joint behavior is clearly smoother and more accurate, even though the actuator still sticks and slips.




Figure 4.16: Mechanically Decoupled Design; Joint-Position Response




Figure 4.17: Mechanically Decoupled Design; Actuator-Position Response

Figures 4.18 and 4.19 show the additional effect of increasing outer-loop bandwidth to the 10 Hz that stability analysis in Section 4.3.1 indicates is possible with the Joint/Actuator design. Joint tracking becomes even more accurate.




Figure 4.18: Joint/Actuator Design; Joint-Position Response




Figure 4.19: Joint/Actuator Design; Actuator-Position Response

There is evidence in these figures that stick-slip behavior in actuator response can lead to limit-cycling, with a frequency that increases with the controller bandwidth. The point of view taken here is that some actuator limit-cycling is acceptable as long as endpoint behavior remains smooth, as in Figures 4.18 and 4.19 above. In some applications this may not be true; actuator wear or power consumption might increase to significant levels. Techniques to suppress the limit cycle such as series compensation (see for instance [32]), or the addition of a deadband to the controller, should then be investigated.

As an additional note which will not be pursued here, dither can be added to the actuator signal to reduce the effects of dry friction^[14,23]. The increased disturbance rejection of the Joint/Actuator controller design at high frequencies would allow a higher value of dither amplitude than the reference design. Of course, dither has the same drawbacks as limit cycling in terms of increased power consumption and wear, so its utility depends on the application.

Overall, these simulations support the conclusion that the mechanical decoupling of a manipulator's joint and actuator nodes can considerably improve smoothness and tracking accuracy in the presence of dry friction. The higher bandwidth permitted by the Joint/Actuator controller gives an additional benefit.

4.3.4 Contact Behavior

A smooth, rapid transition from free motion to constrained motion is highly desirable, but difficult to obtain with a rigid environment. The stiffness controllers considered in this thesis have first-order target dynamics $\left(\frac{1}{d_T s + k_T}\right)$, which would provide the desired smooth transition if:

- i) they are tracked perfectly by the controller, and
- ii) they are applied directly to the interaction force between payload and environment.

Since every controller has a bandwidth limit, the first condition does not hold initially during a collision; response is essentially open-loop. This case will be examined first. At some point the controller does make itself felt, so closed-loop behavior is investigated next.

The second condition above is also not strictly true, because the force sensor is inboard of the payload mass. In fact, in this model the force sensor location is the dividing line between what is considered link mass and payload mass. Inertial payload dynamics can give rise to force oscillations and bouncing during contact even if tracking is perfect.

Due to the discontinuity at contact, linear analysis is difficult to apply. The simulations in the next chapter will explore contact behavior in more detail, but some simple arguments are used here to indicate the roles of drive compliance, bandwidth, and payload mass.

Open-Loop Response

To first investigate *open-loop* contact behavior, a simplified version of the manipulator is considered. Actuator position y is held fixed and a moving environment of stiffness k_E impinges on the link with velocity $-V$. Drive compliance $d_{Ds} + k_D$ allows the endpoint position x to react. The desired result is the time history of F , the interaction force between the manipulator and the environment. From this can be determined the peak collision force and the time of the first “bounce.” The assumption of fixed actuator position during this first bounce should be valid for high-transmission-ratio manipulators, since these have low drive stiffness, high effective actuator mass, and significant actuator friction, all of which reduce the backdriveability of the actuator. A stiff inner position loop on the actuator also contributes to this effect.

The dynamics are given by the following equations:

$$\begin{aligned} (m_L + m_P)\ddot{x} + d_D \dot{x} + k_D x &= -F \\ F &= k_E (x + Vt) \end{aligned} \quad (4.3.20)$$

Expressed in Laplace notation, the interaction force is given by:

$$F(s) = \frac{V k_E [(m_L + m_P)s^2 + d_D s + k_D]}{s^2 [(m_L + m_P)s^2 + d_D s + k_D + k_E]} \quad (4.3.21)$$

Solving this differential equation using the inverse Laplace transform gives:

$$F(t) = \frac{V k_E}{k_E + k_D} \left[k_D t + \frac{2k_E \zeta}{\omega_n} + k_E e^{-\zeta \omega_n t} \left(\frac{1 - 2\zeta^2}{\omega_{dn}} \sin \omega_{dn} t - \frac{2\zeta}{\omega_n} \cos \omega_{dn} t \right) \right] \quad (4.3.22)$$

using the following definitions:

$$\omega_n \equiv \sqrt{\frac{k_D + k_E}{m_L + m_P}}$$

$$\zeta \equiv \frac{d_D}{2\sqrt{(m_L + m_P)(k_D + k_E)}} \quad (4.3.23)$$

$$\omega_{dn} \equiv \omega_n \sqrt{1 - \zeta^2} = \sqrt{\frac{k_E + k_D}{m_L + m_P} - \frac{d_D^2}{4(m_L + m_P)^2}}$$

The oscillatory component decays with the following time constant:

$$\tau = \frac{1}{\zeta \omega_n} \quad (4.3.24)$$

Equation (4.3.22) is valid until contact with the surface is first lost. Since drive damping is typically small ($\zeta \approx .018$ for the values in Table 4.1) its effect during the first bounce can be neglected, giving:

$$F(t) \approx \frac{V k_E}{k_E + k_D} \left(k_D t + \frac{k_E}{\omega_n} \sin \omega_n t \right) \quad (4.3.25)$$

Solving for the first $t > 0$ at which $F(t) = 0$ (the time at which the manipulator loses contact with the surface) requires:

$$\frac{\sin \omega_n t}{\omega_n t} = -\frac{k_D}{k_E} \quad (4.3.26)$$

Although an explicit solution for t is unavailable, it can be seen that as $\frac{k_D}{k_E}$ approaches zero, $t \rightarrow \frac{\pi}{\omega_n}$, and as $\frac{k_D}{k_E}$ increases, $t \rightarrow \frac{3\pi}{2\omega_n}$. For small k_D , the bounce time approaches:

$$t_{\text{bounce}, k_D=0} = \pi \sqrt{\frac{m_L + m_P}{k_E}} \quad (4.3.27)$$

As k_D increases from zero to $k_D = .217 k_E$, this time increases by 29 %. If $k_D > .217 k_E$, there is no solution for bounce time; the manipulator remains in contact with the surface.

The peak value of interaction force reached during the first bounce can be calculated by differentiating eqn. (4.3.25), solving for the time of maximum and evaluating the force at that time:

$$F_{\text{max}} = \frac{V k_E}{k_E + k_D} \left[k_D \cos^{-1} \left(-\frac{k_D}{k_E} \right) + \sqrt{k_E^2 - k_D^2} \right] \sqrt{\frac{m_L + m_P}{k_E + k_D}} \quad (4.3.28)$$

For $k_D < k_E$ the effect of drive compliance is small; as k_D increases from zero to k_E the peak force increases by only 11 %. For $k_D = 0$, eqn. (4.3.28) reduces to the usual value for an object of mass $(m_L + m_P)$ bouncing on a surface of stiffness k_E :

$$F_{\text{max}, k_D=0} = V \sqrt{(m_L + m_P) k_E} \quad (4.3.29)$$

Some conclusions can now be drawn about initial open-loop behavior during contact with a stiff environment ($k_E > k_D$). The qualitative behavior, i.e. whether or not a bounce occurs, is independent of collision velocity, although the magnitude of the forces generated is proportional to it. The response has an oscillatory component which decays with time constant τ , given by eqn. (4.3.24).

The peak force and oscillation frequency are determined primarily by the contact velocity, link + payload mass, and environmental stiffness. Decreasing drive stiffness can reduce peak interaction force, but the effect is slight; the maximum force is well approximated by eqn. (4.3.29). A stiff drive ($k_D > .217 k_E$) can suppress bouncing regardless of contact velocity or payload, although if the drive becomes too stiff the assumption of no actuator motion is invalid and the actuator will simply bounce with the link. The time of the first link bounce (if it occurs) is approximately given by eqn. (4.3.27).

The best prospects for improving open-loop contact behavior through mechanical design remain the generally-recognized principles of minimizing link mass and limiting effective environmental stiffness. Drive compliance has a relatively small effect.

Closed-Loop Response

At some point in the manipulator's response its closed-loop behavior becomes important, as the controller reacts to measured interaction forces. The above analysis indicates that a controller could help damp out the oscillations if its bandwidth exceeds $\zeta\omega_n$, but to reduce the peak force or suppress the initial bounce, the bandwidth must approach ω_{dn} .

The closed-loop contact behavior of the Joint/Actuator controller and the reference design can be compared by the same Laplace ramp-response method used above for the open-loop case, but using the full linear manipulator model and controller equations instead of the simple fixed-actuator version. Equation (3.4.8) is used along with the controller transfer functions (3.4.16) and (4.3.1) to obtain $I(s)$, the closed-loop admittance. The payload and environment are represented by the usual mass, damping, and stiffness parameters m_P , d_E , and k_E , but now has velocity $-V$. The following equations result:

$$x = -I(s) F$$

(4.3.30)

$$F = m_P s^2 x + (d_E s + k_E) \left(x + \frac{V}{s^2} \right)$$

Combining these equations gives:

$$F(s) = \frac{V(d_E s + k_E)}{s^2[1 + I(s)(m_p s^2 + d_E s + k_E)]} \quad (4.3.31)$$

The symbolic solution for time response in this case is so complicated as to be unenlightening, so the response is obtained numerically (using the *MATLAB*TM computer program) for the standard example parameters in Table 4.1, assuming collision velocity $V = 1$ m/sec. Since force response is simply proportional to V , the forces for any other velocity can be obtained by scaling.

Figure 4.20 shows the interaction force $F(t)$ for the initial bounce of the reference design controller, compared with that of the target dynamics. Note that again this linear analysis assumes continuous contact with the environment, so these results are only valid up to the point the manipulator leaves the surface; approximately when $F(t)$ crosses zero at .0227 seconds. This bounce time agrees well with the open-loop estimate from eqn. (4.3.27), which gives .0222 seconds. The peak force during the first bounce is 140.6 N, also agreeing with the calculated open-loop value from eqn. (4.3.29) of 141.4 N.




Figure 4.20: Reference Design; Contact Force Response

Figure 4.21 shows the same time period for the Joint/Actuator controller. Since contact is continuous this trajectory is valid over the whole interval.




Figure 4.21: Joint/Actuator Design; Contact Force Response

The plot interval can be extended to show the decay of the force oscillations:




Figure 4.22: Joint/Actuator Design; Contact Force Response, Extended Interval

The force oscillations from impact decay with a time constant of approximately 0.05 sec, eight times as fast as the open-loop response from eqn. (4.3.24).

It is clear from this comparison that the Joint/Actuator controller design can improve contact behavior over that of the reference design, which essentially gives the open-loop initial response. In this case the peak impact force decreased 14 % and bouncing was suppressed.

These results and the open-loop analysis indicate that the improvement is attributable mostly to the increased overall bandwidth of the Joint/Actuator controller, rather than the reduced drive stiffness. Contact behavior can be improved over the open-loop response only when overall bandwidth (ω_j in the Joint/Actuator case) approaches ω_{dn} (from eqn. 4.3.23), the frequency of the payload + link-environment resonance. Since overall bandwidth in the reference design is limited by the actuator-drive resonance ω_D , contact behavior in that case can be improved over open-loop response only if the environment is more compliant than the drive, or if the link + payload inertia is much greater than the actuator inertia. With the Joint/Actuator design this is not a constraint, and contact behavior can be improved significantly.

It is interesting to note that if for a particular application bounce suppression is more important than minimizing peak force magnitude, link mass should be *increased*. This reduces the frequency of the oscillations, allowing the controller a better chance to damp them out before a bounce occurs.

Effect of Payload Mass

Although the target dynamics in a stiffness controller are first-order in the measured force, any inertia between the force sensor and the environment (i.e. payload mass) can introduce oscillatory behavior into the target trajectory. In the trajectory of a real manipulator this is added to the tracking errors studied above, since it is present even if the target dynamics are tracked perfectly.

To investigate this effect on contact behavior, the previous example was modified by increasing payload mass from .01 kg to 1.0 kg, equal to the link inertia. Note that payload mass is not modeled internally in any controller considered here. The following figure compares the resulting force trajectory of the Joint/Actuator controller to the new target trajectory:




Figure 4.23: Joint/Actuator Design w/Payload; Contact Force Response

Increasing payload mass lowers the frequency of the open-loop force oscillations, permitting the controller to damp them out more effectively. However, there is now noticeable oscillation in the target trajectory. This generates peak forces larger than those due to tracking error in this case; for high payload masses it will limit the best contact behavior achievable with any controller.

The foregoing analysis shows that it is the *inertial* character of the manipulator which limits contact performance, giving rise to high impact forces and bouncing. In impedance-control terms, the bandwidth of the controller implies an equivalent closed-loop link inertia, with a high-bandwidth controller reducing this effective inertia, and low bandwidth giving the open-loop (physical) link inertia. The contact forces due to tracking error are the forces of this equivalent inertia colliding with the environment. The real inertia of the payload must be added to this to determine total inertia, and the total contact forces.

A simple argument indicates that even with an ideal controller, the total effective inertia of the system cannot be reduced below the actual payload inertia: consider a generic manipulator carrying a payload of mass m_p , with a force sensor in between. Assume that the manipulator has a controller which can give it any desired acceleration. To make the manipulator behave as if it has target mass m_T , the desired acceleration in response to external forces is given by $\ddot{x} = -\frac{F}{m_T}$, where F is the compressive force measured by the sensor. For causality, assume that this acceleration takes place some small time Δt after the force measurement, so $\ddot{x}(t + \Delta t) = -\frac{F(t)}{m_T}$. Then in free space, where $F(t) = m_p \ddot{x}(t)$, one obtains:

$$\frac{\ddot{x}(t + \Delta t)}{\ddot{x}(t)} = -\frac{m_p}{m_T} \quad (4.3.32)$$

If $m_T < m_p$, the magnitude of the acceleration diverges to infinity. The controller is unstable if the target mass is less than the payload mass, regardless of how small Δt is. Such a stability limit has been observed in experiments with impedance controllers^[56].

The design recommendation that follows from this is to keep the inertia of the manipulator structure *outboard of the force sensor* to a minimum; below the equivalent closed-loop link inertia. Payload mass obviously depends on the task at hand, but any manipulator structural mass which contributes to it limits the best obtainable contact behavior.

Force-sensing fingers are preferable to a force-sensing wrist, for example. As an analogy, consider the difficulty of picking up an egg between two heavy objects held in the hands. A much more delicate touch is possible when one's force-sensing fingers are directly applied to the egg.

4.4 Conclusions: Performance Comparison Results

This section summarizes the results of this chapter's performance comparison between the Joint/Actuator controller and the reference design, for the example parameters in Table 4.1. The following table shows ten measures related to the performance criteria defined in Section 3.1.2:

Parameter	Reference Design Value	Joint/Actuator Design Value	Relative Improvement
Open-Loop Transfer Function Magnitude @ 100 Hz	-2.47 db	-2.73 db	.26 db
Phase Margin	16.4°	51.1°	312 %
Usable Bandwidth	4 Hz	20 Hz	500 %
Steady-State Stiffness Error	- 37.5 %	0	∞
Steady-State Actuator Force Disturbance Rejection	-44.58 db	-78.15 db	33.57 db
Steady-State Actuator Position Disturbance Rejection	59.58 db	39.96 db	19.62 db
Transmission Force Disturbance Rejection by Inner Loop	0 db	20.83 db	20.83 db
Static Force Dynamic Range	*	*	20–30 db
Peak Force in 1 m/s Impact	140.6 N	120.4 N	14 %
Bouncing at Contact	yes	no	Bouncing suppressed

Table 4.3: Performance Parameter Comparison

The first two parameters above are related to stability robustness. Both controllers are stable for the operating conditions of the example; these values give an indication of the sensitivity of this stability to small changes in the manipulator or environment. Since it has been assumed that the model may be inaccurate above 100 Hz, the open-loop transfer function gain at 100 Hz gives an indication of sensitivity to unmodeled dynamics above that frequency, with more attenuation preferred. The values for the two controllers are virtually

* Value depends on magnitude and nature of disturbances.

the same, with a slight advantage going to the Joint/Actuator design. Both offer attenuation at all frequencies above 100 Hz. Phase margin, the second parameter, is the amount of additional phase lag at the gain crossover frequency required to bring the system to the verge of instability, and indicates relative stability. A standard rule of thumb is that phase margin should be between 30° and 60° for satisfactory performance^[38]. The Joint/Actuator design provides a significantly better phase margin.

Usable bandwidth for the reference controller is limited by a resonance peak in the response due to drive compliance; for frequencies above that peak, closed-loop behavior differs greatly from the target dynamics. For the Joint/Actuator controller usable bandwidth is limited to the point at which response begins to roll off from the target value. A useful increase in bandwidth is obtainable with the new controller.

Stiffness error gives an indication of the relative accuracy of the tracking of the target dynamics within the usable bandwidth; for ease of comparison it is evaluated at steady state (zero frequency). Drive compliance leads to significant error with the reference design; there is none with the new controller.

The disturbance rejection figures reflect the relative smoothness of force control. For the above table, force and position disturbance rejections are evaluated at zero frequency (their relative magnitudes remain approximately constant throughout the usable bandwidth). Note that position disturbance rejection is the magnitude of a position-to-force transfer function, and is not a dimensionless ratio, but the improvement tabulated above is a ratio of these magnitudes and is dimensionless. Force and position disturbance rejection are both greatly improved with the Joint/Actuator design.

As Section 4.3.3 showed, the position disturbance rejection advantage is reflected in superior tracking of a constant-velocity trajectory with friction and stiction. Stick-slip behavior at the actuator was effectively filtered out of the endpoint trajectory with the new design.

The rejection of transmission force disturbances by the inner loop alone is also compared. This represents the error in applied joint force due to friction, etc. arising in the transmission. Errors in endpoint force would be further reduced by the outer loop in each case, but this comparison points out just the difference due to inner loop structure. In the Joint/Actuator design, the inner loop reduces these forces by a factor of $(N + 1)$, where the drive has been decoupled by decreasing its stiffness by a factor of N . No attenuation is provided by the reference design; all rejection of transmission force disturbances must come from the outer loop.

Actual values for force dynamic range depend on the type of disturbances (force or position) present and their magnitudes, so numerical values are not shown. The relative

improvement of the Joint/Actuator controller over the reference design lies in the range of 20–30 db, however.

Peak impact force is an indication of the quality of contact behavior, and should be small. The tabulated value is for the parameter values of Table 4.1, and assumes a 1 m/s collision; the force scales proportionally with velocity. In this case, the Joint/Actuator controller suppressed bouncing from the surface at contact; the reference design did not.

By every criterion used the Joint/Actuator controller outperforms the reference design, in many cases by a wide margin.

Chapter 5: Simulation and Experiments

This chapter completes the analysis by examining time-domain behavior and including nonlinear effects. Numerical simulation and experimentation with a single-joint manipulator testbed are used for this purpose.

The simulation is used to confirm the predicted behavior of the controllers described in Chapters 3 and 4, to discover the effects of modeling error, and to investigate techniques for improving the practicality of the Joint/Actuator design by reducing actuator force requirements. Experiments are performed to validate the manipulator dynamic model used throughout this thesis, and to directly demonstrate the stability advantages of the new controller.

5.1 Simulation

The methods used for simulation are discussed in the first section. Next, the comparison example of Chapter 4 is investigated, first with a low-speed trajectory and then with a high-speed trajectory. It is seen that the predicted improvements in stability, smoothness, bandwidth, and contact behavior indeed materialize in the simulation. As expected, however, when full advantage is taken of the improved bandwidth capability of the Joint/Actuator controller, the required actuator forces are much higher.

After a look at the effects of modeling error in the fourth section, the emphasis turns to ways of reducing actuator force levels. Saturation and trajectory filtering are examined in the next two sections, which succeed in reducing the peak force level to 22 % of its initial value without degrading performance.

The following section concentrates on the stability advantages of the new controller. It is shown that for low target stiffnesses, superior performance can be obtained with the Joint/Actuator design with *no* penalty in actuator force.

Conclusions drawn from the simulation work are summarized in Section 5.1.8.

5.1.1 Description of Simulation

Numerical simulation allows the inclusion of a number of effects present in a real manipulator which are not treated in continuous-time linear analysis. The simulation used here was written in *C* and run on a Macintosh SE/30 computer.

Digital implementations of the reference design control law and the Joint/Actuator control law were derived using procedures described in reference [6]. Zero-order-hold equivalents are used for the filter dynamics $\left(\frac{g_{D0}s + g_{P0}}{m_L s^2 + g_{D0}s + g_{P0}}\right)$ in the reference design and for the target dynamics $\left(\frac{1}{\hat{d}_T s + \hat{k}_T}\right)$ in both designs. Experimentation with the simulator showed that the drive model $\left(\frac{1}{\hat{d}_D s + \hat{k}_D}\right)$ in the Joint/Actuator controller could be adequately approximated by simply $\frac{1}{\hat{k}_D}$ for the small amounts of drive damping assumed.

This approximation gave lower control forces and lower sensitivity to modeling errors than the zero-order-hold equivalent of the original expression. The drive damping estimate \hat{d}_D is still used in the feedforward of drive compression force.

Finite (adjustable) resolution is assumed for measurements of position and force, as well as the control input U . The simulation includes actuator saturation, contact discontinuity, and friction nonlinearities. Provision is also made for introducing modeling error into the controllers.

The dynamic equations are integrated using a 4th-order Runge-Kutta scheme with a time step of .0001 second. The maximum controller sampling rate simulated is 1000 Hz, giving at least 10 dynamics-updates per sampling period.

In addition to the simulated controller behavior, the program calculates the target trajectory (in both position and interaction force) which would result if the manipulator perfectly tracked the commanded stiffness.

Since the algorithms being simulated control stiffness, their input is a reference position trajectory. The standard trajectory used begins in free space at $x = -1$ cm, with the environment consisting of a wall at $x = 0$. At time $t = 0$, the manipulator is at rest with no position error. The reference point then moves at constant velocity of 2 cm/sec for 1 second, placing it 1 cm past the surface of the wall. It remains stationary for 1 second, then moves back out at -4 cm/sec for 0.5 second, placing it back at $x = -1$ cm. The simulation continues for a further 1 second with the reference point stationary.

This trajectory gives a look at the response to sharp changes in the reference trajectory in free space and while in contact with the environment, the response to impact and to leaving a surface, the tracking of a constant reference velocity, and the settling to a constant desired force while in contact, and to a constant position in free space.

These velocities are typical of an approach to contact that might be used in space applications, where the components to be handled may be delicate and speed is secondary in importance to avoiding damage. In an industrial setting much higher speeds are desirable, so an example using a high-velocity trajectory is simulated in Section 5.1.3. Results are qualitatively similar, so in all other examples the low-speed trajectory above is used.

5.1.2 Controller Comparison Example

The comparison example analyzed in Chapter 4 is simulated first. The parameter values from Table 4.1 are used. The following table shows additional simulation parameters used for this example:

Parameter	Value
Sampling Rate	1000 Hz
Actuator Saturation Limit	∞
Coulomb Friction F_C	20 N
Stiction ΔF_{stict}	50 N
Viscous Friction d_A	10 Ns/m
Estimated Viscous Friction \hat{d}_A	8 Ns/m
Estimated Actuator Inertia \hat{m}_A	1.1 kg
Estimated Link Inertia \hat{m}_L	1.1 kg
Estimated Drive Stiffness \hat{k}_D	110 N/m
Estimated Drive Damping \hat{d}_D	0.45 Ns/m
X Sensor lsb Value	.00016 m
Y Sensor lsb Value	.000041 m
F Sensor lsb Value	.0122 N
Control Input lsb Value	.0122 N

Table 5.1: Controller Comparison Simulation Parameters

The estimated manipulator parameters include an intentional modeling error of 10 %, except for actuator viscous friction, which has 20 % error. The “lsb values” shown represent the resolution of the sensors simulated. No actuator saturation limit is imposed for this example.

Figure 5.1 shows the manipulator endpoint (X) trajectory results for the reference design controller. Evidence of stick-slip behavior can be seen in the initial free-space motion. Several bounces occur upon contact with the surface at $x = 0$. Sticking also delays response when motion reverses. The effects of the lightly-damped drive resonance can be clearly seen after $t = 2.5$ seconds, as it settles to a fixed commanded position. The oscillations occur at about 5 Hz, which is the drive resonance frequency ω_D .

Figure 5.2 shows the corresponding trajectory for the Joint/Actuator controller. No bouncing occurs at contact and the tracking is very good.




Figure 5.1: Reference Design; Simulated Position Trajectory




Figure 5.2: Joint/Actuator Design; Simulated Position Trajectory

A more sensitive indication of tracking performance, and one more relevant to force control, is given by the trajectory of interaction force F . Figures 5.3 and 5.4 show this trajectory for the reference design and the Joint/Actuator design, respectively:




Figure 5.3: Reference Design; Simulated Force Trajectory




Figure 5.4: Joint/Actuator Design; Simulated Force Trajectory

The reference design shows a great deal of oscillation in the force response. It also settles toward an incorrect steady-state value (38 % error), due to drive compliance.

The Joint/Actuator design tracks the target trajectory quite well, with some force overshoot at sharp transitions (such as contact) due to limited bandwidth. In the steady-force phase between $t = 1$ and $t = 2$ seconds, there is a low-amplitude limit cycle which further simulation shows is due to stiction; the force remains within .05 N (0.9 %) of the target value.

The main drawback of the new controller is that to take full advantage of the improved bandwidth capability (giving the excellent tracking shown above), much higher actuator forces are required. This is true whenever bandwidth is increased, but it is exaggerated in this case by the drive decoupling, which increases the amplitude of actuator motion required. Figure 5.5 shows the control input U as a function of time for the reference design:




Figure 5.5: Reference Design; Simulated Control Input Trajectory

The overall control bandwidth in the reference design case is roughly equal to the drive resonance frequency. In the Joint/Actuator case, the bandwidth is nearly 7 times the drive frequency. Equation (4.1.1) predicts approximately 23 times the peak actuator force of the reference case will be required for the Joint/Actuator case to track the same trajectory,

just considering the additional actuator acceleration needed. Dissipation by actuator friction will increase this even more. As Figure 5.6 shows, this is borne out by the simulation; the peak-force ratio is approximately 70 in this example.




Figure 5.6: Joint/Actuator Design; Simulated Control Input Trajectory

In the Joint/Actuator case the actuator force required is very frequency-dependent. In the steady-force region between $t = 1$ and $t = 2$ seconds, the force is negligible compared to that required during sharp transitions. The *average* actuator force required is not 70 times that of the reference case. In this example, the ratio of average absolute values of actuator force is 10.4.

Comparing Figures 5.5 and 5.6 also indicates there is more high-frequency content in the control signals for the Joint/Actuator case. The author's speculation is that this is due to the conditional stability of the new controller. Equation (4.3.8) shows that the actuator-loop bandwidth must be above a certain value to give stability; this indicates that the full actuator bandwidth is being utilized to stabilize the system, not just improve disturbance rejection and tracking. Some high-frequency control activity is then expected, even when motion is smooth.

It is possible that this high-frequency control content will cause problems in practical situations. Some types of motors react poorly to high-frequency control signals--commutator brushes, if present, may arc and erode prematurely. Motor inductance,

neglected here, may prevent the motor from responding as modeled at high frequencies. Experiments described in the latter half of this chapter show that the predicted performance is, in fact, realizable with hardware. No degradation attributable to this effect is observed, since the discrepancies that occur between simulation and experiment are approximately the same for both types of controller. Nevertheless, care must be taken in manipulator design to recognize these as potential problems.

This comparison in this section confirms the advantages in smoothness, bandwidth, and contact behavior of the Joint/Actuator design over the reference design predicted by the analysis in Chapter 4. As expected, the principal cost of this approach is the need for higher actuator force capability. Ways of reducing this cost are investigated in later sections discussing the effects of saturation and filtering of the reference trajectory.

5.1.3 A High-Speed Example

This example repeats the previous comparison, but using a high-speed trajectory. The trajectory begins in free space at $x = -10$ cm. At time $t = 0$, the manipulator is at rest with no position error. The reference point then moves at constant velocity of 50 cm/sec for 0.4 second, placing it 10 cm past the surface of the wall. It remains stationary for 1.6 seconds, then moves back out at -100 cm/sec for 0.2 second, placing it back at $x = -10$ cm. The simulation continues for a further 1.3 seconds with the reference point stationary. The controller parameters from tables 4.1 and 5.1 are used.

Figure 5.7 shows the position trajectory obtained with the reference design controller, and Figure 5.8 shows that of the Joint/Actuator design.




Figure 5.7: Reference Design; Simulated Position Trajectory




Figure 5.8: Joint/Actuator Design; Simulated Position Trajectory

The interaction force trajectories are shown in Figures 5.9 and 5.10:




Figure 5.9: Reference Design; Simulated Force Trajectory




Figure 5.10: Joint/Actuator Design; Simulated Force Trajectory

Comparison with the previous example shows that the results are qualitatively the same, although numerical values of forces and position errors scaled up with the velocities. Higher reference velocities do not significantly change the character of performance achieved by the two control schemes.

5.1.4 Effects of Modeling Error

The Joint/Actuator controller design depends on a model of the drive dynamics, given by the parameters \hat{d}_D and \hat{k}_D . If the controller's performance or stability were very sensitive to errors in these parameter estimates, it would be of limited value in the real world.

To test this, some simulations were performed with large errors in the modeled drive parameters. It was found that errors of 50 % or more in drive damping are tolerated with minor degradation of performance. Figure 5.11 shows the simulated force trajectory obtained with +50 % error in drive stiffness \hat{k}_D and -50 % error in drive damping \hat{d}_D :




Figure 5.11: Simulated Force Trajectory, +50 % \hat{k}_D , -50 % \hat{d}_D Modeling Error

Figure 5.12 shows the results with +50 % error in \hat{d}_D :




Figure 5.12: Simulated Force Trajectory, +50 % \hat{k}_D , +50 % \hat{d}_D Modeling Error

The trajectories are essentially the same and reasonably approximate the target. Errors of $\pm 50\%$ in drive damping thus have very little effect in this example. The error in drive stiffness \hat{k}_D causes an error in steady-state force applied. Steady-state force error due to stiffness modeling error can be predicted from the transfer functions in (4.3.1), and is given by:

$$\frac{F_{\text{error}}}{F} = (\hat{k}_D - g_{Pa}) \left(1 - \frac{\hat{k}_D}{k_D} \right) \left(\frac{k_E k_T}{k_E + k_T} \right) \left(\frac{1}{g_{Pa} g_{Pj}} \right) \quad (5.1.1)$$

which can usually be approximated by:

$$\frac{F_{\text{error}}}{F} \approx \frac{k_T}{g_{Pj}} \left(\frac{\hat{k}_D}{k_D} - 1 \right) \quad (5.1.2)$$

In this case a steady-state force error of 4.5 % results from a 50 % stiffness overestimate.

The effects of underestimating drive stiffness \hat{k}_D are somewhat more severe. Figure 5.13 shows the force trajectory with -30 % error in stiffness and -50 % in damping:




Figure 5.13: Simulated Force Trajectory, -30 % \hat{k}_D , -50 % \hat{d}_D Modeling Error

Comparison with Figure 5.11 shows increased oscillation (ringing) in the force response, as well as the expected steady-state error. Control input levels also increased: the average absolute control level is 4.8 times that of the trajectory of Figure 5.11.

These results indicate that it is better to overestimate drive stiffness than to underestimate it, and that large (30-50 %) errors in the estimate are tolerable. Accuracy of the drive damping estimate is not critical, with ± 50 % error producing no noticeable effect.

5.1.5 Effects of Actuator Saturation

Real actuators have finite force capabilities; if the control law calls for more force, the output saturates. Sweet and Good^[48] refer to actuator saturation as “the most significant nonlinearity” in the control of manipulators. If occasional actuator saturation does not degrade performance excessively, smaller actuators can be used, saving mass and power.

The stability criterion (4.3.8) for the Joint/Actuator controller gives an indication of the effects of saturation on its performance. Saturation of the control signal of the inner loop (the actuator force) in effect lowers the inner-loop gains for high error values. By

eqn. (4.3.8) the inner loop is conditionally stable, i.e. lowering inner-loop bandwidth reduces stability. Therefore, excessive saturation is expected to lead to instability.

Simulation indicates that for the example of Table 5.1, the maximum control input called for is 5328 N. Imposing a saturation limit of ± 3000 N has no significant effect on performance, but a limit of ± 2500 N causes instability. This instability occurs at $t = 2$ seconds, when the reference trajectory instantaneously accelerates from rest to -4 cm/sec. Once the error is large enough to cause instability, it continues even after the reference trajectory stops moving. This is illustrated by Figure 5.14. In the graph the vertical range is kept small to show some detail in the early part of the trajectory; as a result the bouncing which occurs after $t = 2$ seconds goes off-scale.




Figure 5.14: Simulated Position Trajectory; ± 2500 N Saturation Limit

It is interesting to note from eqn. (4.3.8) that, unlike the inner loop, the outer (joint) loop is not conditionally stable; decreasing outer-loop bandwidth *improves* stability. Saturation of the *joint-force* command should therefore not lead to instability, but (more benignly) degrade bandwidth and tracking quality. Referring to equation (4.2.1), the Joint/Actuator control law can be written:

$$u = (\hat{d}_D s + \hat{k}_D)(y - x) + (g_{Da} s + g_{Pa}) \left[\left(\frac{1}{\hat{d}_D s + \hat{k}_D} \right) F_{JD} + x - y \right] \quad (5.1.3)$$

$$\text{where } F_{JD} \equiv F + (g_{Dj} s + g_{Pj}) \left(x_{\text{ref}} - \frac{F}{d_T s + k_T} - x \right)$$

This equation makes explicit the joint-force command F_{JD} , which consists of feedforward of the sensed interaction force F , and a term applying the outer-loop gains to the joint position error. F_{JD} is calculated in the controller algorithm as an intermediate step, and it is easy to impose a saturation limit on it in the software. It should also be noted in passing that feedforward compensation of disturbance forces acting on the joint axis (such as friction) can easily be added to F_{JD} if a disturbance model is available, although this is not pursued here.

Imposing a saturation limit on the joint loop could provide an effective safety feature, to reduce the demands on the inner loop and limit the divergence which could occur when the inner loop saturates. This effect is illustrated in Figure 5.15, which uses the same parameters as the simulation in Figure 5.14, except that a 8 N saturation limit is imposed on the commanded joint force F_{JD} .

Divergence still occurs at $t = 2$ seconds when the actuator initially saturates, but now the manipulator is able to recover stability and resume accurate tracking. The force trajectory shown in Figure 5.16 shows that tracking quality is not significantly degraded by the limit on commanded joint force, as long as enough force is allowed to properly interact with the environment during stable operation. The errors at $t > 2$ are due to the actuator force limit.




Figure 5.15: Simulated Position Trajectory; using Joint Force Saturation Limit




Figure 5.16: Simulated Force Trajectory; using Joint Force Saturation Limit

To summarize the effects of saturation: in this example, limiting the actuator force to 50 % of its peak value resulted in instability; higher limits had no discernible effect. The consequences of actuator saturation can be ameliorated, without degrading tracking performance, by an appropriately-chosen limit on commanded joint force.

5.1.6 Effects of Filtering the Reference Trajectory

The study of saturation indicates that the highest actuator forces, and therefore the closest approach to saturation-induced instability, occur at sharp transitions in the reference trajectory. This suggests that filtering unnecessary high frequencies from the reference trajectory could reduce actuator force requirements.

To test this, simulations were run with the reference trajectory passed through a first-order low-pass filter with adjustable corner frequency. The best results were found with the corner frequency equal to the joint-loop bandwidth, at 10 Hz. Frequency content above this value is not trackable anyway, so performance is not impaired by the filtering.

With this filter, the actuator saturation limit could be reduced from ± 3000 N to ± 1200 N without triggering instability, and the average absolute actuator force required decreased by 10 %.

Figure 5.17 shows the tracking obtained with this filter and saturation limit. Comparison with Figure 5.4 shows an improvement in smoothness over the unsaturated, unfiltered trajectory. Figure 5.18 illustrates the control input trajectory. Comparison with Figure 5.6 shows a 78 % reduction in peak actuator force.




Figure 5.17: Simulated Force Trajectory; ± 1200 N Saturation Limit, Filtered X_{ref}




Figure 5.18: Simulated Control Trajectory; ± 1200 N Saturation Limit, Filtered X_{ref}

5.1.7 Low-Stiffness Example

The simulations in Section 5.1.2 confirm the predicted advantages of the Joint/Actuator controller in bandwidth, smoothness, and contact behavior. The other principal advantage expected is in stability range. Lower commanded stiffness should be attainable when the Joint/Actuator design is used.

To test this, simulations were performed with a lower target stiffness. The simulation parameters are listed in the following table:

Parameter	Value
Target damping d_T	10 Ns/m
Target stiffness k_T	25 N/m
Environment damping d_E	2 Ns/m
Environment stiffness k_E	5000 N/m
Sampling Rate	1000 Hz
Coulomb Friction F_C	2.0 N
Stiction ΔF_{stict}	5.0 N
Viscous Friction d_A	1.0 Ns/m
Estimated Viscous Friction \hat{d}_A	.8 Ns/m
Estimated Actuator Inertia \hat{m}_A	1.1 kg
Estimated Link Inertia \hat{m}_L	1.1 kg
Estimated Drive Stiffness \hat{k}_D	110 N/m
Estimated Drive Damping \hat{d}_D	0.45 Ns/m
X Sensor lsb Value	.00016 m
Y Sensor lsb Value	.000041 m
F Sensor lsb Value	.0122 N
Control Input lsb Value	.0122 N

Table 5.2: Low-Stiffness Simulation Parameters

Other values remain as in Table 4.1. The X_{ref} trajectory used is the low-speed one of Section 5.1.2, but is filtered as described in Section 5.1.6.

For the reference design controller a wide range of gains and filter bandwidths were tried, without achieving good results. If bandwidth is set too high, instability occurs upon

contact; if set lower, tracking is poor. These cases are illustrated by the position trajectories in Figures 5.19 and 5.20, and the force trajectory in Figure 5.21:




Figure 5.19: Reference Design; Simulated Position Trajectory, $\omega_o = 2$ Hz




Figure 5.20: Reference Design; Simulated Position Trajectory, $\omega_o = 0.1$ Hz




Figure 5.21: Reference Design; Simulated Force Trajectory, $\omega_o = 0.1$ Hz

In each of these cases, the X_{ref} trajectory was filtered with a corner frequency of 10 Hz, and the inner-loop bandwidth ω_i was set to 50 Hz. In Figure 5.19, with an outer-loop bandwidth of 2 Hz, instability set in upon contact. This resulted in a series of bounces of increasing amplitude, until the reference trajectory left the surface.

For the simulation of Figures 5.20 and 5.21 outer-loop bandwidth was lowered to 0.1 Hz. This value gave the best performance of any values tested, and required a peak actuator force of 13.2 N. Bouncing was reduced, but position and force tracking were very poor. The reference design controller was unable to produce practically useful performance at this low target stiffness.

The Joint/Actuator controller was simulated with a range of bandwidths as well. Figure 5.22 shows the position response with an inner-loop bandwidth of 50 Hz and outer-loop bandwidth of 10 Hz. Figure 5.23 shows the force trajectory for the same case.




Figure 5.22: Joint/Actuator Design; Simulated Position Trajectory, $\omega_j = 10$ Hz




Figure 5.23: Joint/Actuator Design; Simulated Force Trajectory, $\omega_j = 10$ Hz

Tracking of both position and force was very good in this case. Actuator forces were very high, with a peak of 4663 N and an absolute average of 363 N. It seemed likely that a lower bandwidth would still yield acceptable performance and should reduce actuator force requirements, so further simulations were performed.

Reducing the joint-loop bandwidth to 5 Hz resulted in the force trajectory of Figure 5.24. It is still quite good, although open-loop behavior dominates the initial contact response, causing a brief bounce and some oscillation. Actuator forces were 1412 N peak and 88.5 N average, a reduction of approximately a factor of 4.




Figure 5.24: Joint/Actuator Design; Simulated Force Trajectory, $\omega_j = 5$ Hz

To pursue this further, simulations were run with a joint-loop bandwidth of 1 Hz. The X_{ref} trajectory was filtered with a corner frequency of 1 Hz as well in this case, and actuator saturation was imposed at a level of ± 13 N, resulting in the position trajectory of Figure 5.25 and the force trajectory of Figure 5.26. The average actuator force required was 4.52 N. Compared to the $\omega_j = 10$ Hz case, the peak actuator force was therefore reduced by a factor of 358, and the average by a factor of 80.

Three bounces occur at contact, but within 1 second the force error settles to a small value, as one would expect with an overall bandwidth of 1 Hz. The free-space position tracking is fairly good. This level of performance is useful in many applications, and is

much better than that shown in Figures 5.20 and 5.21 for the reference design. It should be noted that this was achieved with *less* actuator force.




Figure 5.25: Joint/Actuator Design; Simulated Position Trajectory, $\omega_j = 1$ Hz




Figure 5.26: Joint/Actuator Design; Simulated Force Trajectory, $\omega_j = 1$ Hz

The results of this low-stiffness example show that the Joint/Actuator controller has the predicted advantage in stability range. At this target stiffness the new design significantly outperformed the reference design, demonstrating much more usable performance even when restricted to the same actuator force levels. At higher actuator force levels, extremely accurate tracking could be obtained.

5.1.8 Conclusions

Simulations confirmed many of the qualitative predictions of the previous analysis and revealed additional information on modeling-error sensitivity and actuator force requirements.

The Joint/Actuator controller proved superior in all four of the performance criteria used: stability, bandwidth, smoothness, and contact behavior. Although it depends on a drive model, it can tolerate large errors (50 %) in the estimated parameters with little effect on performance. The high tolerance for errors in drive damping \hat{d}_D is particularly useful, since this is a difficult parameter to measure accurately and may change with operating conditions such as temperature and drive tension.

The major disadvantage of this approach is the higher actuator force requirement, due to a control bandwidth above the drive natural frequency. This is exacerbated by the fact that, although peak force requirements can be cut about 50 % with saturation, the response to excessive saturation is instability. The following points should be noted, however:

- In many situations, trading higher actuator force for better performance is advantageous. With the reference design, performance could not be improved to this level at all.
- Average actuator forces are not increased as much as peak forces.
- Filtering unneeded high frequencies from X_{ref} significantly reduces force requirements, and in fact improves smoothness.
- Imposing a saturation limit on joint force commands is easily done in controller software and reduces the consequences of actuator saturation instability.
- Even when bandwidth is lowered to use actuator forces comparable to the reference design, superior performance is possible.

The specific design recommendations which follow from the simulation work are to tend to overestimate k_D for the drive model rather than underestimate it, include a joint-force saturation limit, and filter unneeded high frequencies from the reference trajectory X_{ref} . If possible, actuators should be selected that can briefly exert peak forces much higher than the sustainable average, and they should be tolerant of high-frequency control signals.

5.2 Experiments

The linear analysis and the nonlinear simulations show a number of advantages for the new controller and joint design. All of these results were based on a simplified model of real mechanisms, described in Section 2.6. Although this model was suggested by studies of actual manipulators^[16,48,51], the work here pushes the model in new directions. To see whether the predicted improvements carry over into the real world of hardware, experiments with the new design were performed on a single-joint testbed.

The primary purpose of the experiments is to validate the simulation procedure used above. The predicted and achieved trajectories of position and force are compared for several operating points. It is shown that the simulation accurately predicts the onset of contact instability as the target stiffness is lowered. Oscillation frequencies and settling times are close to the expected values.

A second goal of the experiments is to directly demonstrate the stability advantage of the new approach. This is done by comparing the performance of the Joint/Actuator design to the reference design for the same set of target dynamics and environment. These results agree with the simulation as well, and the stability improvement is clear.

5.2.1 Description of Hardware

The experimental hardware is depicted in the figure below:




Figure 5.27: Experimental Apparatus, Top View

Actuation is provided by a permanent-magnet DC motor (Pittman model # GM 9414E729). A 500-line incremental optical encoder (the “Y” sensor) is mounted on one end of the motor shaft, and a gearbox with ratio of 38.3 is mounted on the other end. To the output shaft of the gearbox is attached a 1-inch diameter capstan for a cable drive.

The joint is located approximately 1 meter away. The actuator capstan is connected to the joint capstan by 1/16”-diameter braided stainless steel aircraft control cable. The cable includes a lightweight turnbuckle for tensioning, and interchangeably incorporates either two drive-decoupling springs (used with the Joint/Actuator controller) or rigid aluminum pieces of the same length and mass (used with the reference design controller). The ends of the cable are anchored to the joint capstan by set-screws. Four and one-half turns of cable are wrapped around the actuator capstan, and affixed in the middle to the capstan by another set-screw. Cable tension is set at 30-35 N.

Two angular-contact ball bearings support the joint shaft. The lower end of the shaft is coupled to a 5000-line incremental optical encoder (the “X” sensor). A 4-inch diameter cable drive capstan is attached to the middle of the shaft, and the link is attached to the upper end of the shaft. The joint has approximately 180° of travel before encountering

rubber bump stops. The transmission ratio requires 76.6 revolutions of the motor shaft to produce this range of motion.

The link is a .5-inch diameter horizontal steel rod approximately 1 meter in length, with the joint axis intersecting vertically at the center. A counterweight is clamped to one end of the link and the "wrist" force sensor attached to the other. The radius from joint axis to force-sensing tip is 52 cm. Link motion takes place in a horizontal plane, so gravity effects are negligible. The first bending mode of the link has an observed natural frequency of about 25 Hz.

The force sensor is an aluminum alloy (6061-T6) beam instrumented with a strain-gage bridge to measure interaction forces in the (circumferential) direction of link motion. The required flexibility of the force sensor limits the effective environmental stiffness to a maximum of 2970 N/m. The force signal is boosted by an instrumentation amplifier located at the wrist, and relayed to a 12-bit A/D converter.

The X and Y encoders are read by circuits using 16-bit counters. The circuit for the X encoder uses quadrature to augment position resolution by a factor of four. Sufficient resolution is available already from the Y encoder, due to the large transmission ratio. Velocity measurements are obtained with circuits which time the interval between successive encoder pulses. These intervals are inverted in software to give velocities. A consequence of this method is that velocity resolution becomes finer at low velocities. The clock rates used are selected to give at most $\pm 2\%$ quantization error at the highest expected endpoint velocity of 4 m/sec.

Data is gathered and control torque calculated by an IBM PC AT personal computer with a 16-bit parallel interface to the sensor circuitry. The software is written in Microsoft C (Version 5.0), and runs at a cycle rate of 535 Hz. A listing of the controller software appears in Appendix C.

A 12-bit D/A converter drives an op-amp (*Apex PA-03*) in a circuit which converts the voltage command to a current output, producing the desired torque from the actuator.

The limitations of this particular hardware introduced two effects which have not been included in previous simulations. The first is a constraint on output *voltage* of the actuator drive circuit, which imposes a torque limit which is velocity-dependent (due to varying back-emf of the DC motor). This saturation limit is given by:

$$10.6 (4.09 \text{ m/s} - \dot{y}) > U > 10.6 (-4.09 \text{ m/s} - \dot{y}) \quad (5.2.1)$$

where U is control force (in N), and \dot{y} is actuator velocity (in m/s). This limit exists in addition to the current-constrained saturation limit at a fixed value of ± 27.1 N.

The second effect is related to the sampling of the somewhat noisy force signal F . An analog filter was applied to this signal before sampling to eliminate aliasing and reduce the effects of electrical noise. The peak-to-peak noise amplitude of the unfiltered signal was approximately 50 lsb units, corresponding to 0.66 N of force. A first-order RC low-pass filter with bandwidth of 6 Hz reduced the peak-to-peak noise value to approximately 5 lsb units, or 0.07 N of force. Since the maximum overall control bandwidth used in the experiments is 1.5 Hz, this filtering should not degrade performance.

These two effects were included in the simulations presented in the next section for comparison with experiments.

An effect not included in the simulations is backlash in the actuator or transmission. In the implementation of the Joint/Actuator controller, the effect of backlash is to add a deadband to the joint force command. In other words, initial movement of the actuator with respect to the joint produces no joint force until the backlash is taken up. This can be compensated for by including a bias force in the drive model: positive bias force if the desired force command is positive, negative bias if the command is negative.

This type of compensation was used in the experiment controller, with a bias force magnitude of 0.13 N. Rather than add backlash *and* its compensation into the simulation, both were left out.

The most likely sources of remaining discrepancies between the simulation and experiment are errors in the modeled parameters, sensor noise, joint-axis friction, link bending, and actuator cogging; the simulation includes none of these effects.

The following table summarizes the parameters of the hardware, as seen from the endpoint of the force sensor:

Parameter	Value
Payload Mass m_P	.052 kg
Link Mass m_L	.739 kg
Actuator Mass m_A	.206 kg
Drive Stiffness k_D	2870 N/m
Drive Damping d_D	3.58 Ns/m
Drive Stiffness k_D (decoupled)	192.5 N/m
Drive Damping d_D (decoupled)	.24 Ns/m
Environmental Stiffness k_E	2970 N/m
Environmental Damping d_E	2.5 Ns/m
Coulomb Friction Force F_C	.77 N
Stiction Force ΔF_{stict}	4.0 N
Viscous Friction d_A	.64 Ns/m
X Sensor lsb Value	.156 mm
Y Sensor lsb Value	.0427 mm
Velocity lsb Value (after conversion)	1.248 mm/s
F Sensor lsb Value	.024 N
Control Input lsb Value	.01323 N
Sampling Rate	535 Hz
Control Force Saturation Level	± 27.1 N
Actuator-Drive Frequency ω_D	18.8 Hz
ω_D (decoupled)	4.87 Hz
Link-Environment Frequency ω_E	9.75 Hz
Payload-Environment Frequency ω_P	38.0 Hz

Table 5.3: Experimental Apparatus Parameters

The values for physical parameters above were calculated from measurements made on the assembled hardware. Calibrated weights were used to provide known forces for the stiffness measurements. Inertias and damping values were calculated from curvefits to trajectories measured in response to a constant applied force.

5.2.2 Results

Due to the limited torque available from the actuator, very high bandwidths could not be explored with the present apparatus. So, rather than verifying the predicted bandwidth improvements of the Joint/Actuator approach, the experiments were selected to confirm its stability advantages at low target stiffnesses. This can be done without large actuator forces, as shown by the low-stiffness example of Section 5.1.7.

The results of four experimental runs are presented. The parameters used for each run are shown below:

Run #	Controller Type	Drive	Inner-Loop Bandwidth	Outer-Loop Bandwidth	k_T	d_T
1	Joint/Actuator	Decoupled	$\omega_a = 30$ Hz	$\omega_j = 1.5$ Hz	50 N/m	30 Ns/m
2	Joint/Actuator	Decoupled	$\omega_a = 30$ Hz	$\omega_j = 1.5$ Hz	15 N/m	9 Ns/m
3	Reference	Stiff	$\omega_i = 30$ Hz	$\omega_o = 1.5$ Hz	50 N/m	30 Ns/m
4	Reference	Stiff	$\omega_i = 30$ Hz	$\omega_o = 0.1$ Hz	50 N/m	30 Ns/m

Table 5.4: Experimental Parameters

The achievable sampling rate of 535 Hz limited the inner-loop bandwidth to 30 Hz in each case. The maximum outer-loop bandwidth of 1.5 Hz for the Joint/Actuator controller was set by saturation limits.

Run 1 was chosen to demonstrate stable contact with a rigid environment and relatively soft target stiffness, using the new controller design. Run 2 was chosen to demonstrate behavior on the edge of stability: the target values were found by trial-and-error to be ones which produce bouncing from the surface which neither grows in magnitude nor damps out (when the target trajectory should be producing constant force). Comparison with the simulation for this borderline case provides a sensitive test of the simulator's accuracy in predicting contact stability.

Run 3 directly compares the reference design controller with the Joint/Actuator controller of Run 1. It shows instability upon contact. Run 4 reduces the outer loop bandwidth in an attempt to improve stability. Bouncing still occurs, and tracking is degraded.

The experimental trajectory is the low-velocity one used in the simulations of Section 5.1. The “wall” at $X = 0$ in this case is solid aluminum; much more rigid than the arm’s force sensor. It is therefore the force sensor which sets the effective environmental compliance.

The force trajectories shown here are slightly different than those in Section 5.1. Since the purpose is to compare with the experiment, the simulated forces plotted are not the interaction forces between the tip of the manipulator and the environment as before, but those that would be measured by the force sensor itself. The difference is the inclusion of the inertial force due to acceleration of the payload mass (outboard of the force sensor), and also simulated sensor quantization and filtering. This makes the simulated forces directly comparable to the experimentally measured ones.

In each case the experimental or simulated trajectory is plotted with the “target” trajectory for reference, representing the trajectory of an ideal manipulator tracking the target dynamics.




Figure 5.28: Experimental Position Trajectory, Run 1




Figure 5.29: Simulated Position Trajectory, Run 1




Figure 5.30: Experimental Force Trajectory, Run 1




Figure 5.31: Simulated Force Trajectory, Run 1

The experimental results differ from the simulation in having larger free-space position tracking errors, convergence to a slightly incorrect final position, and a notably “rounder” force peak.

The final convergence error is attributable to drift in the force sensor calibration during the run, leading to an incorrect value for the desired equilibrium position. This effect can be seen to some extent in all of the runs.

Although this is speculation, some stiction at the joint axis could produce the observed rounding of the force peak, by distorting the relationship between joint position and endpoint force when the tip is nearly stationary. It would also delay motion when direction reverses, contributing to the tracking errors seen when the tip comes off of the surface.

The general behavior of the experiment is close to that of the simulation, however. Oscillation frequencies and settling times are quite similar. Some oscillation at approximately 12 Hz is observable in both force trajectories, and oscillation at about 3 Hz is present in both free-space position trajectories. The steady-state force values achieved are close to the same, and are approached at similar times. Qualitatively, the overall level of stability, or damping, in the experimental results matches that of the simulation.




Figure 5.32: Experimental Position Trajectory, Run 2




Figure 5.33: Simulated Position Trajectory, Run 2




Figure 5.34: Experimental Force Trajectory, Run 2




Figure 5.35: Simulated Force Trajectory, Run 2

Once again, in free space the experimental tracking error is larger than that of the simulation. The position amplitude of the bounces is higher by a factor of 2-3.

The force response of the experiment displays some ringing at a frequency of about 25 Hz, which is not present in the simulation. This is due to excitation of the first bending mode of the link, which is not modeled in the simulation.

The overall response is quite similar. The ramping pattern of the force bounces during the constant-velocity phase and the approximately constant-amplitude bounces that follow are a close match in amplitude and frequency. The bounces occur at about 7 Hz in the simulation and 6 Hz in the experiment, reaching a maximum amplitude of about 1.3 N in each.

This comparison indicates a particularly good prediction of the stability of the experiment. Slightly stiffer target dynamics produce decaying bounces, and slightly softer dynamics produce divergence. Both the simulation and experiment display bouncing of approximately constant amplitude.




Figure 5.36: Experimental Position Trajectory, Run 3




Figure 5.37: Simulated Position Trajectory, Run 3




Figure 5.38: Experimental Force Trajectory, Run 3




Figure 5.39: Simulated Force Trajectory, Run 3

The position amplitude of the rebound is again higher by a factor of 2-3 in the experiment than in the simulation, but the shape of the bounces is similar. The frequency of the bouncing itself is about 3 Hz in the simulation and 2 Hz in the experiment. A higher-

frequency force oscillation superposed on the bounces occurs at about 16 Hz in both the simulation and experiment. The force amplitude of the bouncing reaches a peak of about 3 N in each.

As the simulation and previous analysis predict, for these target dynamics the performance of the reference design controller is inferior to the Joint/Actuator performance of Run 1. Contact with the surface results in bouncing instability in this case, where stable contact was achieved before. Although it is difficult to make quantitative estimates of disturbance rejection from this data, evidence of stick-slip behavior can be seen in the position response of both the experiment and simulation for the reference design, during the constant-velocity phases of the trajectory.




Figure 5.40: Experimental Position Trajectory, Run 4




Figure 5.41: Simulated Position Trajectory, Run 4




Figure 5.42: Experimental Force Trajectory, Run 4

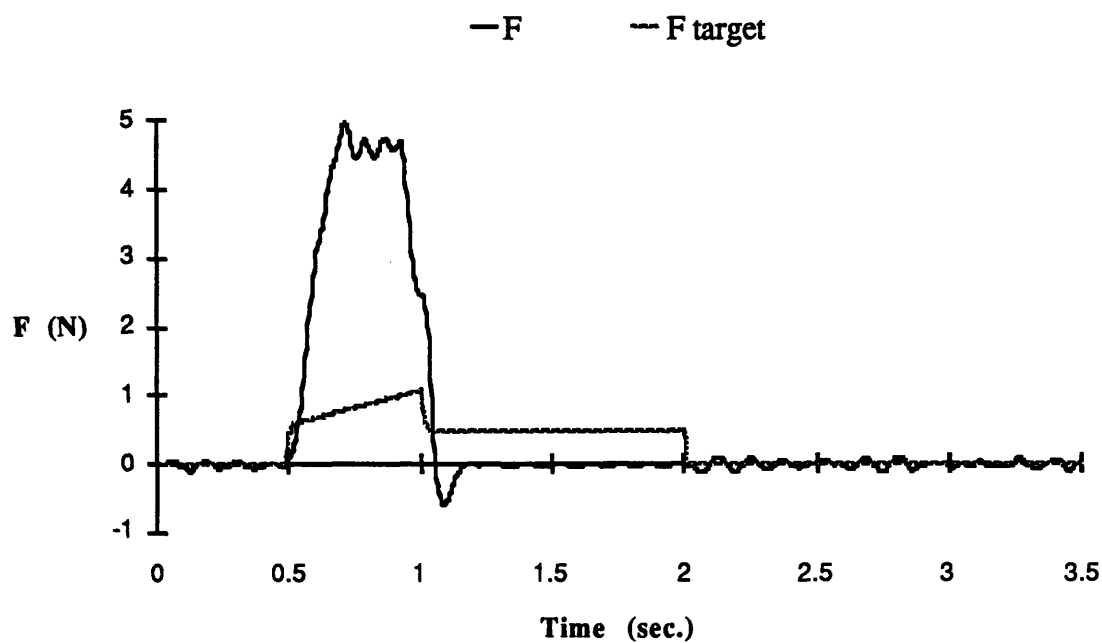


Figure 5.43: Simulated Force Trajectory, Run 4

This case exhibits particularly close correspondence between the simulation and experiment. This might be due to overall lower high-frequency content of the trajectories, because of the reduced control bandwidth.

The response consists of a single bounce, reaching a peak force of about 5 N in each case. Thus, reducing the outer-loop bandwidth to the low value of 0.1 Hz did not succeed in producing stable contact with the reference design controller. Position and force tracking are significantly degraded by the reduced bandwidth. Note that once again, stick-slip behavior can be seen in the position response.

Less formal experiments performed with this hardware indicate that the superior stability of the Joint/Actuator approach extends to softer target dynamics with softer environments. Specifically, when the force sensor is pushed with a person's finger or hand, the environmental stiffness is quite low. The Joint/Actuator controller can produce smooth, stable response to a finger-push with target dynamics of $k_T = 1$ N/m and $d_T = 2$ Ns/m, while the reference design responds with (violent) instability.

5.2.3 Conclusions

In the important qualitative areas of smoothness and stability, the simulation correctly predicts the behavior of actual hardware using the control algorithms studied. Quantitative predictions of settling time, oscillation frequency, and force amplitude are also quite good.

The largest errors occur in the prediction of free-space position tracking performance, with both types of controller. The higher errors seen in the experiments could be caused by joint-axis disturbance forces (e.g. friction and stiction) which are unmodeled in the simulation. In a practical application a manipulator designed to take full advantage of the Joint/Actuator controller design would use a higher joint-loop bandwidth than the current experiment, giving better rejection of joint-axis disturbance forces. Also, if a model of joint-axis disturbance forces is available, feedforward compensation in the joint force command F_{JD} can be used to reduce errors.

The experiments also confirm the results of the low-stiffness simulation of Section 5.1.7. They show that the Joint/Actuator design can improve stability for low target stiffnesses even without increasing actuator force requirements.

These results indicate that the simple manipulator dynamic model adopted in Section 2.6 is adequate to explain the performance of actual hardware using force control, validating the analysis of Chapter 4. The experiments demonstrate that real performance advantages can be obtained with the Joint/Actuator design approach.

- Joint/Actuator control design (introduced here) combined with mechanical drive decoupling yields significant performance benefits over current designs in the areas of stability, bandwidth, smoothness, and contact behavior.
- The advantages of the Joint/Actuator approach over force-controlling inner-loop methods are better rejection of transmission disturbances, the use of a colocated sensor which allows high gains, and the fact that it is position-based allowing feedforward friction compensation at both the joint and actuator.
- The principal cost of this approach is an increase in actuator force requirements, particularly when full advantage is taken of the bandwidth improvement.
- Although the Joint/Actuator controller depends on a drive model, simulation shows that modeling errors of +50 % -30 % in drive stiffness and ± 50 % in drive damping are tolerated with minor loss of performance.
- Simulation also indicates that actuator force limits and reference trajectory filtering can reduce peak actuator forces by a factor of five.
- The response of the Joint/Actuator controller to excessive actuator saturation is instability; recovery from this instability is facilitated by the use of *joint force* limits in software.
- Stability advantages of the Joint/Actuator design at low commanded stiffnesses are demonstrable in simulation and with hardware, and do not require large actuator forces.
- The dynamic model and simulation procedure used throughout the thesis are validated by experimental data.

This thesis addresses the problem of improving the force-control performance of manipulators with transmission dynamics. Analysis, simulation, and experiments reveal the shortcomings of current designs and show that the new Joint/Actuator design can provide a significant improvement, in the typical case where drive compliance is the performance-limiting factor. The analyses of controller stability tradeoffs, contact behavior, and saturation limits also contribute new insight into the mechanics of manipulators and force control algorithms in general.

6.2 Design Recommendations

Joint/Actuator manipulator design broadens the range of force-control performance that can be attained with a non-direct drive manipulator. It also relaxes the need to keep the manipulator drive mechanically stiff, giving the engineer more latitude in transmission design. However, joint position sensors must be added, and provision made for the drive-decoupling compliance. For high-bandwidth operation, large actuators will be required.

This section outlines a suggested design procedure for a force-controlling manipulator, based on the ideas and analyses presented in this thesis. Design is always an iterative process. The approach described here starts with performance goals and calculates what is required to achieve them. The goals may then have to be revised and the process repeated. A popular alternative is to begin with the constraints on cost, size, and mass and then design for the maximum performance that can be obtained, revising the constraints later if necessary. Although it may seem that this cost-centered method would produce a more efficient design, the reverse is often true. Since the first approach forces the designer to define the required performance up front, a more focused and useful design often results. Also, the second approach never comes in *under* budget.

The range of intended tasks must first be carefully defined. The designer begins with an idea of the required bandwidth, static force capability, payload range, minimum desired stiffness, maximum desired stiffness, expected maximum collision velocity, maximum allowable contact force, estimated disturbance forces and required smoothness of operation.

Maximum required stiffness determines the maximum environmental spring constant k_E that must be dealt with in the following way. Installing a compliant covering on the manipulator grippers allows the designer to set an upper limit on the effective k_E seen by the manipulator. As shown in Chapter 4, reducing k_E lowers contact forces, reduces bouncing, and improves stability. However, a compliant covering also limits the maximum effective manipulator stiffness as seen by the payload and environment. The compliant covering should therefore be used to reduce k_E to its minimum acceptable value, i.e. the maximum required manipulator stiffness.

The minimum required stiffness, combined with knowledge of the payload range and the requirement that commanded behavior be overdamped, sets a minimum commanded manipulator damping value d_T .

The required control bandwidth sets a minimum value for the joint-loop bandwidth ω_j . Another minimum value for ω_j is set by the contact requirements. With k_E , payload,

and maximum collision velocity now fixed, if maximum allowable contact force is exceeded by the resulting open-loop behavior (eqn. 4.3.29), the only remedy is to increase ω_j to at least the damped natural frequency of contact ω_{dn} (from eqn. 4.3.23).

From these values for k_E , minimum d_T , and minimum ω_j , the stability criterion eqn. (4.3.8) leads to a minimum value for the actuator-loop bandwidth ω_a .

The required position disturbance rejection combined with the already-set environment parameters give a maximum value for the drive stiffness k_D , to be obtained by adding a linear elastic element to the transmission. A certain minimum amount of drive decoupling is also required for proper controller operation: to ensure sufficiently accurate modeling of drive compliance, and to obtain sufficient resolution in the controller's differencing of actuator and joint position to give a force estimate. The latter limit depends on the resolution of the position sensors used, and the former can be checked by calibration.

With minimum values for the control gains and a maximum drive stiffness now set, the force disturbance rejection can be computed (eqn. 4.3.14) and applied to the expected disturbance forces. If the resulting trajectory smoothness does not meet requirements, two remedies are available: increasing bandwidths further or decreasing drive stiffness. Both increase power consumption and actuator force requirements. Increasing bandwidths may require faster sampling rates and will increase sensitivity to unmodeled dynamics by increasing the gain crossover frequency of the open-loop transfer function (eqn. 4.3.6). On the other hand, decreasing drive stiffness will increase force and power requirements with no benefit in speed of response (eqn. 4.1.1). This tradeoff must be made with a detailed knowledge of the costs. As mentioned above, iteration of the design is generally required.

Some specific hardware recommendations can be made. Since the controller can cope with a large amount of drive compliance and the actuators are likely to be massive, it is particularly advantageous to place all of the actuators on the manipulator base. Transmissions like the tension-element drives used by Townsend^[51] should be used to convey actuator forces to the joints.

The next recommendation is to locate the speed-reducing part of the transmission (e.g. gearbox) at the actuator axis rather than the joint axis, with the decoupling spring between the transmission and the joint. This has two benefits: it minimizes inertia by moving the gearbox mass off of the moving links, and it maximizes rejection of force disturbances arising in the transmission (eqn. 4.2.5).

The third hardware recommendation is to reduce structural mass outboard of the force sensor to a minimum. As described in Section 4.3.4, any such inertia limits the best

obtainable contact behavior. For a delicate touch, force-sensing skin is preferable to force-sensing fingers, which are preferable to a force-sensing wrist, etc.

Finally, if actuator limit cycles prove to be a problem, the designer should recognize that increasing drive compliance will reduce their effect on endpoint motion. Of course, standard treatments such as adding a controller deadband are still available.

Some of these design recommendations conflict with those suggested for other force controller designs. Townsend^[51], for instance, recommends increasing k_D , lowering k_E until $\omega_E = \omega_D$, and locating the speed-reduction at the joint. This is sound advice for the standard type of controller, because its bandwidth and disturbance rejection are limited by instability due to drive dynamics. Increasing the natural frequency of drive dynamics allows higher gains.

In the Joint/Actuator controller it is the inner-loop dynamics which provide an analogous limit for the outer-loop gains, so in this case stability is improved by increasing inner-loop bandwidth. Drive stiffness no longer affects stability and is chosen from other considerations; primarily disturbance rejection.

6.3 Recommendations for Further Work

In the examples used in this thesis, actuator inertia was set equal to link inertia (except in the experiment, where measured values differed somewhat). This corresponds to choosing the “impedance-matching” transmission ratio, which gives the highest joint acceleration for a given actuator torque capability, in the rigid-drive case.

When the drive is decoupled, however, this ratio is no longer “optimal” by the same criterion. The proportion of actuator torque which is conveyed to the joint is now frequency-dependent. Finding the “optimal” transmission ratio will require a more thorough system analysis, trading off actuator torque costs with desired bandwidth. Since this tradeoff is very task-specific it was not pursued at this stage of research. The topic deserves further investigation, since lower actuator inertia would reduce required torques significantly.

Another area for study would be the addition of feedforward to the joint loop, to compensate for known disturbances. In standard controllers there is no way to directly compensate for joint friction, since the sensing and control take place at the actuator. In the Joint/Actuator design, however, joint position is directly sensed and joint force commands are explicitly formulated. A model of friction could easily be included, if available. This potential benefit remains unexploited in this thesis.

Manipulators designed to be very lightweight or carry massive payloads may have low-frequency link bending modes. These non-colocated modes limit performance in the same way as drive compliance. The Joint/Actuator design described here does not cope with link bending any better than a standard design, because link dynamics occur outboard of the joint sensor. The Joint/Actuator design could be extended, however, to use endpoint sensing instead of joint sensing. Small-amplitude link bending would then be lumped in with the deflections due to drive compliance. The deflection of the (softer) drive-decoupling spring would dominate this effect as well as the original transmission compliance, and the controller design should improve stability.

Finally, the principles developed here can be generalized to other classes of mechanism. In addition to enhancing stability in force control, the Joint/Actuator controller approach combined with mechanical drive decoupling is a powerful method for improving rejection of mechanical disturbances. It can be thought of as a mechanical analog of electrical noise-reduction systems (e.g. *Dolby™*).

The figure below shows a system with a source of additive noise between the input and output. One approach to noise reduction is to pre-filter the input signal, boosting its level at frequencies where noise is significant, then filter the output signal with an inverse of the pre-filter. The input waveform is restored and the noise is attenuated. This practice is common in the recording industry, where the processes of analog recording, duplication and playback add high-frequency noise to a signal.

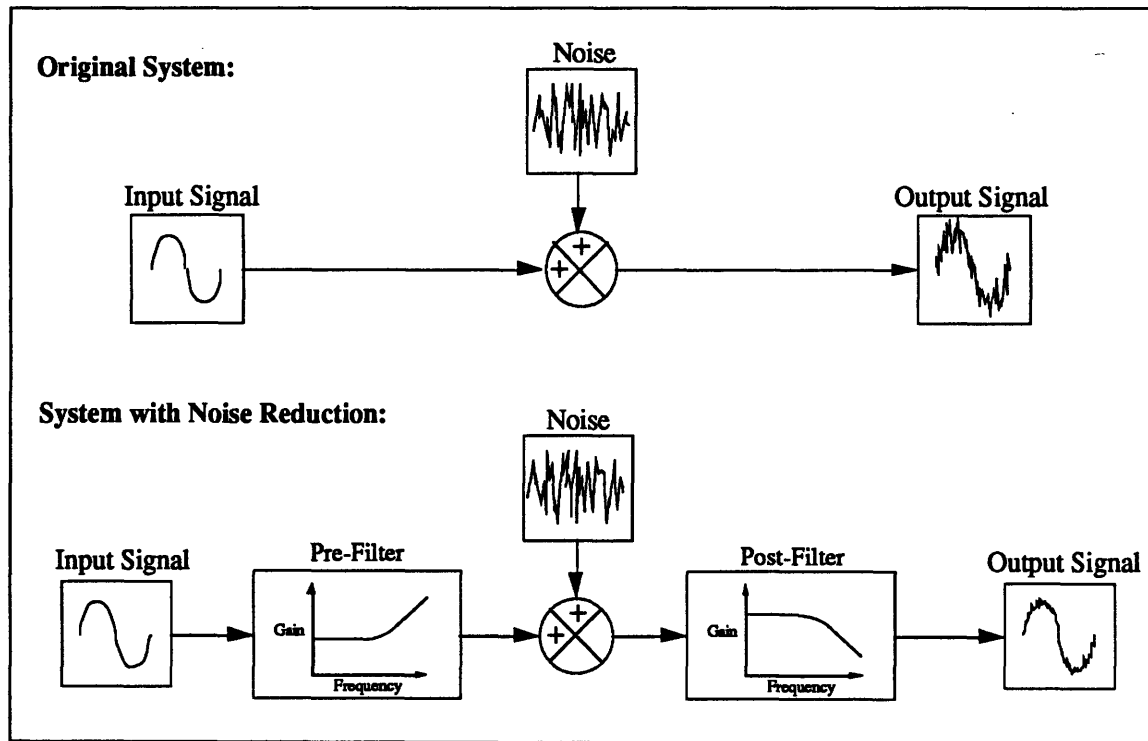


Figure 6.1: Noise-Reduction Analogy

In the Joint/Actuator manipulator analogy, the post-filter is provided mechanically by the drive-decoupling compliance, and the pre-filter is the inverse drive model in the controller software. The “noise” consists of actuator disturbance forces such as friction, backlash, and cogging. As this thesis demonstrates, the effect of these disturbances on the output can be reduced significantly.

A wide range of mechanisms could conceivably benefit from this technique. Although a simple spring is used to decouple the manipulator drive in this thesis, more sophisticated mechanical filters could certainly be designed to attenuate specific disturbances.

Appendix A: Manipulator Mass Scaling Model

It is assumed that each actuator by itself can produce the required tip acceleration A with all other joints fixed. This gives:

$$\begin{aligned}\dot{\omega}_1 &= A/3L \\ \dot{\omega}_2 &= A/2L \\ \dot{\omega}_3 &= A/L\end{aligned}\tag{A.1}$$

Each link is assumed to be a uniform beam of length L . With the geometry shown in Figure 2.1, this gives:

Direct-Drive

$$\begin{aligned}J_{T1} = J_1 &= (9 M_P + \frac{19}{3} M_{L3} + \frac{7}{3} M_{L2} + \frac{1}{3} M_{L1} + 4 M_3 + M_2) L^2 \\ J_{T2} = J_2 &= (4 M_P + \frac{7}{3} M_{L3} + \frac{1}{3} M_{L2} + M_3) L^2 \\ J_{T3} = J_3 &= (M_P + \frac{1}{3} M_{L3}) L^2\end{aligned}\tag{A.2a}$$

Geared

$$\begin{aligned}J_1 &= [9 M_P + \frac{19}{3} M_{L3} + \frac{7}{3} M_{L2} + \frac{1}{3} M_{L1} + 1.44 (M_3 + M_{G3}) + .04 (M_2 + M_{G2})] L^2 \\ J_2 &= [4 M_P + \frac{7}{3} M_{L3} + \frac{1}{3} M_{L2} + .04 (M_3 + M_{G3})] L^2 \\ J_3 &= (M_P + \frac{1}{3} M_{L3}) L^2\end{aligned}\tag{A.2b}$$

For the geared manipulator, actuator rotor inertia is included (neglected in this approximation for the direct-drive case, since it is most significant with large ratios):

$$\begin{aligned} J_{T1} &= J_1/N_1^2 + J_{A1} \\ J_{T2} &= J_2/N_2^2 + J_{A2} \\ J_{T3} &= J_3/N_3^2 + J_{A3} \end{aligned} \quad (\text{A.3})$$

Some assumptions need to be made to determine link masses. Design of a manipulator link depends on many factors, but when designing for minimum mass the most important constraints are required stiffness, length and maximum girth. For simplicity, the links here are assumed to be thin-walled tubes of fixed diameter (presumably determined by geometric constraints of the intended tasks). The thickness of the walls can be varied to change bending stiffness. This gives stiffness proportional to mass for each link.

As a rational way of relating stiffness requirements for the different links, it is assumed here that the worst-case deflection of the payload due to link bending is to be the same for all links, when imparting to the payload the required acceleration. Actual deflections would depend on the materials used and link diameters. By taking M_{L3} (the mass of the third link) as an independent parameter, the deflections and geometry need not be determined. The masses of the other links are set by this mass and the requirement of equal payload deflection under load. M_{L3} can be considered to set the level of stiffness required. These assumptions give the following relationship between link masses:

$$\begin{aligned} \frac{M_{L1}}{M_{L3}} &= \frac{J_1}{J_3} \\ \frac{M_{L2}}{M_{L3}} &= \frac{J_2}{J_3} \end{aligned} \quad (\text{A.4})$$

Combining equations (A.2) and (A.4) gives:

Direct-Drive

$$M_{L1} = \frac{M_{L3}}{M_P} (9 M_P + \frac{19}{3} M_{L3} + \frac{7}{3} M_{L2} + 4 M_3 + M_2) \quad (A.5a)$$

$$M_{L2} = \frac{M_{L3}}{M_P} (4 M_P + \frac{7}{3} M_{L3} + M_3)$$

G geared

$$M_{L1} = \frac{M_{L3}}{M_P} [9 M_P + \frac{19}{3} M_{L3} + \frac{7}{3} M_{L2} + 1.44 (M_3 + M_{G3}) + .04 (M_2 + M_{G2})] \quad (A.5b)$$

$$M_{L2} = \frac{M_{L3}}{M_P} [4 M_P + \frac{7}{3} M_{L3} + .04 (M_3 + M_{G3})]$$

Using Euler's Law $T = J \dot{\omega}_i$; and letting $K = \frac{M_P}{M_{L3}} + \frac{1}{3}$:

$$T_1 = J_1 \dot{\omega}_1 = \frac{1}{3} A L K M_{L1}$$

$$T_2 = J_2 \dot{\omega}_2 = \frac{1}{2} A L K M_{L2} \quad (A.6)$$

$$T_3 = J_3 \dot{\omega}_3 = A L K M_{L3}$$

For the geared case the torques the motors must supply are given by:

$$T_{M1} = J_{T1} N_1 \dot{\omega}_1 = \frac{A}{3} \left(\frac{L}{N_1} K M_{L1} + \frac{N_1}{L} J_{A1} \right)$$

$$T_{M2} = J_{T2} N_2 \dot{\omega}_2 = \frac{A}{2} \left(\frac{L}{N_2} K M_{L2} + \frac{N_2}{L} J_{A2} \right) \quad (A.7)$$

$$T_{M3} = J_{T3} N_3 \dot{\omega}_3 = A \left(\frac{L}{N_3} K M_{L3} + \frac{N_3}{L} J_{A3} \right)$$

The remaining unknowns are motor mass, and (in the geared case) gearbox mass, transmission ratio, and motor rotor inertia.

Scaling laws are proposed in reference [17], based on a simple analysis of the physics involved. The following laws are adopted here:

$$J_{Ai} = \alpha_A (T_{Mi})^{\beta_A} \quad (\text{A.8})$$

$$M_i = \alpha_M (T_{Mi})^{\beta_M} \quad (\text{A.9})$$

$$M_{Gi} = \alpha_G (T_i)^{\beta_G} \quad (\text{A.10})$$

Relationships (A.8) and (A.9) follow from arguments in [17]. The dependence of gearbox mass primarily on output torque T_i in (A.10) is deduced empirically from product data^[59].

Values for the exponents β_A and β_M can be obtained on theoretical grounds^[17]. For the purposes of this model, all values were taken from curvefits to product data^[58,59]. These numbers apply specifically to DC brushless motors, a type suited to space applications. Reasonable agreement exists between the theoretical and empirical values:

<u>Parameter</u>	<u>Empirical</u>	<u>Theoretical</u>
α_A	$1.11 \times 10^{-4} \text{ kg m}^2$	
β_A	1.33	1.22
α_M	3.2 kg	
β_M	.694	.731
α_G	.15 kg	
β_G	.6	

Table A.1: Inertial Scaling Parameters

(Note: These values assume torque is measured in Nm)

The direct-drive case can now be solved. T_3 is calculated from equation (A.6). This gives M_3 through equation (A.9). M_{L2} is then calculated using equation (A.5a). Equation (A.6) gives T_2 , and so on.

The total mass is given by:

Direct-Drive

$$M_T = M_1 + M_2 + M_3 + M_{L1} + M_{L2} + M_{L3} \quad (\text{A.11a})$$

Geared

$$M_T = M_1 + M_2 + M_3 + M_{G1} + M_{G2} + M_{G3} + M_{L1} + M_{L2} + M_{L3} \quad (\text{A.11b})$$

The geared case still contains unknowns: the transmission ratios N_i . Since mass is the “cost” function, the transmission ratios are chosen to minimize mass.

From equations (A.6) and (A.10), M_{Gi} does not depend on N_i . From equation (A.5), M_{Li} does not depend on N_i . Therefore the only mass dependence on transmission ratio is that of M_i . Mass is a monotonic function of motor torque by equation (A.9). Therefore minimizing mass is equivalent to minimizing motor torque T_{Mi} .

The following equations follow from (A.7) and (A.8):

$$\begin{aligned} T_{M1} &= \frac{A}{3} \left(\frac{L}{N_1} KM_{L1} + \frac{N_1}{L} \alpha_A (T_{M1})^{\beta_A} \right) \\ T_{M2} &= \frac{A}{2} \left(\frac{L}{N_2} KM_{L2} + \frac{N_2}{L} \alpha_A (T_{M2})^{\beta_A} \right) \\ T_{M3} &= A \left(\frac{L}{N_3} KM_{L3} + \frac{N_3}{L} \alpha_A (T_{M3})^{\beta_A} \right) \end{aligned} \quad (\text{A.12})$$

To minimize T_{Mi} with respect to N_i , the partial derivatives $\frac{\partial T_{Mi}}{\partial N_i}$ are set to zero.

This gives:

$$N_i^2 = \frac{KM_{Li} L^2}{\alpha_A (T_{Mi})^{\beta_A}} \quad (\text{A.13})$$

Combining equations (A.12) and (A.13):

$$\begin{aligned} T_{M1} &= \frac{2ALKM_{L1}}{3N_1} \\ T_{M2} &= \frac{ALKM_{L2}}{N_2} \\ T_{M3} &= \frac{2ALKM_{L3}}{N_3} \end{aligned} \quad (A.14)$$

In other words, the optimization sets the actuator inertia J_{Ai} equal to $\frac{J_i}{N_i^2}$, the arm inertia as seen from the motor. This is a familiar result, giving the maximum joint acceleration for a given actuator torque. N_i is known as the “impedance-matching” transmission ratio.

Equations (A.13) and (A.14) give:

$$\begin{aligned} N_1 &= \left[\frac{2ALKM_{L1}}{3} \right] \left(\frac{1 - \beta_A}{2 - \beta_A} \right) \left[\frac{3L}{2A \alpha_A} \right] \left(\frac{1}{2 - \beta_A} \right) \\ N_2 &= \left[ALKM_{L2} \right] \left(\frac{1 - \beta_A}{2 - \beta_A} \right) \left[\frac{L}{A \alpha_A} \right] \left(\frac{1}{2 - \beta_A} \right) \\ N_3 &= \left[2ALKM_{L3} \right] \left(\frac{1 - \beta_A}{2 - \beta_A} \right) \left[\frac{L}{2A \alpha_A} \right] \left(\frac{1}{2 - \beta_A} \right) \end{aligned} \quad (A.15)$$

Now equation (A.15) can be solved for N_3 . Equation (A.14) then gives T_{M3} , which gives M_3 by equation (A.9). T_3 is obtained from equation (A.6), giving M_{G3} by equation (A.10). Now equation (A.5b) can be solved for M_{L2} . Equation (A.15) is then solved for N_2 , and so on.

In this way all of the dependent parameters are calculated, finally giving the total mass by equation (A.11).

Appendix B: Analysis of Backlash

A nonlinear control analysis was undertaken to investigate the effects of backlash on a serial-link rotary-joint manipulator. The following section describes the manipulator model and control law assumed. Since these dynamics are nonlinear even without backlash, appeal is made to LaSalle's Theorem to prove global stability. Then backlash is added and a search for potential limit cycles made using Sinusoidal-Input Describing Functions (SIDF's). These results are then checked by numerical simulation of the nonlinear dynamics.

The advantage of the SIDF approach is that it yields symbolic equations for the amplitude and frequency of limit cycles, giving more insight than simulation results alone. The simulations allow assessment of the accuracy of the SIDF method for particular cases.

B.1 Manipulator Model

The following simplified manipulator model is used:




Figure B.1: 2-Link Manipulator Model

This is a two-link manipulator with rotational joints having parallel axes. Motion is therefore confined to a plane. Assume that $L_1 = L_2 = L$ and that the links are thin rods of equal mass M , distributed evenly over their length. No gravity effects are considered.

The model was chosen to be as simple as possible while retaining the nonlinear character of common multi-DOF manipulators. Derivation of the dynamical equations follows standard Lagrangian procedure^[30].

The kinetic energy associated with translation of each link is:

$$E_1 = \frac{1}{8} M L^2 \dot{\theta}_1^2 \quad (B.1)$$

$$E_2 = \frac{1}{2} M L^2 \left[\dot{\theta}_1^2 + \cos(\theta_2) \dot{\theta}_1 (\dot{\theta}_1 + \dot{\theta}_2) + \frac{1}{4} (\dot{\theta}_1 + \dot{\theta}_2)^2 \right]$$

The rotational kinetic energy of each link is given by:

$$R_1 = \frac{1}{24} M L^2 \dot{\theta}_1^2 \quad (B.2)$$

$$R_2 = \frac{1}{24} M L^2 (\dot{\theta}_1 + \dot{\theta}_2)^2$$

Combining these equations one obtains the total kinetic energy of the system, which is also the Lagrangian \mathcal{L} , since there is no potential energy. To simplify the expressions, normalize by setting $\frac{1}{2} M L^2 = 1$:

$$\mathcal{L} = \left[\frac{5}{3} + \cos(\theta_2) \right] \dot{\theta}_1^2 + \left[\frac{2}{3} + \cos(\theta_2) \right] \dot{\theta}_1 \dot{\theta}_2 + \frac{1}{3} \dot{\theta}_2^2 \quad (B.3)$$

Using Lagrange's equation $\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{\theta}_i} \right) - \frac{\partial \mathcal{L}}{\partial \theta_i} = T_i$, one obtains:

$$T_1 = \left[\frac{10}{3} + 2 \cos(\theta_2) \right] \ddot{\theta}_1 + \left[\frac{2}{3} + \cos(\theta_2) \right] \ddot{\theta}_2 - \sin(\theta_2) \dot{\theta}_2 (2 \dot{\theta}_1 + \dot{\theta}_2) \quad (B.4)$$

$$T_2 = \left[\frac{2}{3} + \cos(\theta_2) \right] \ddot{\theta}_1 + \frac{2}{3} \ddot{\theta}_2 + \sin(\theta_2) \dot{\theta}_1^2$$

where T_i is the torque applied to the i^{th} joint. These equations provide the dynamical description of the plant.

B.2 Control Law

A simple linear control law is assumed, using position and velocity feedback on each joint. A fixed set point of $\theta_1 = \theta_2 = 0$ is used, so the joint torques are given by:

$$T_1 = k_1 \theta_1 + k_3 \dot{\theta}_1 \quad (\text{B.5})$$

$$T_2 = k_2 \theta_2 + k_4 \dot{\theta}_2$$

Some performance criteria will be used to set the gains. The system can be roughly approximated by setting the second-order velocity-dependent terms to zero and considering each joint separately. The linearized dynamics become:

$$\frac{16}{3} \ddot{\theta}_1 = T_1 = k_1 \theta_1 + k_3 \dot{\theta}_1 \quad (\text{B.6})$$

$$\frac{2}{3} \ddot{\theta}_2 = T_2 = k_2 \theta_2 + k_4 \dot{\theta}_2$$

These second-order linear equations define a natural frequency and damping ratio for each joint:

$$\omega_{n1} = \sqrt{-\frac{3}{16} k_1}$$

$$\omega_{n2} = \sqrt{-\frac{3}{2} k_2}$$

(B.7)

$$\zeta_1 = -\frac{k_3}{2} \sqrt{-\frac{3}{16} k_1}$$

$$\zeta_2 = -\frac{k_4}{2} \sqrt{-\frac{3}{2} k_2}$$

It seems reasonable to equate the natural frequencies (and damping ratios) for the two joints, which gives:

$$k_2 = \frac{k_1}{8}$$

$$k_4 = \frac{k_3}{8}$$
(B.8)

Later in the analysis, the set of gains will be specified by the controller bandwidth $\omega_c = \omega_{n1} = \omega_{n2}$ and damping ratio ζ of the linearized system (B.6). The actual gains can be calculated from equations (B.7) and (B.8) if desired.

B.3 Stability Analysis

In this section, the second method of Lyapunov will be applied to the system described by equations (B.4) and (B.5) to prove that its behavior is uniformly, asymptotically, stable in the large (UASIL), for appropriately chosen gains. A similar derivation has appeared in [49].

LaSalle's Theorem states^[37]:

For a free dynamic system $\dot{\underline{x}} = \underline{f}(\underline{x}, t)$ where $\underline{f}(\underline{0}, t) = \underline{0}$ and $\underline{f}(\underline{x}, t) = \underline{f}(\underline{x}, t + \tau)$, for all $t \geq 0$: If a scalar function $V(\underline{x}, t) = V(\underline{x}, t + \tau)$ is defined such that

- (i) $V(\underline{0}, t) = 0$
- (ii) $V(\underline{x}, t)$ is positive definite, i.e. there exists a continuous nondecreasing function α such that $\alpha(0) = 0$ and $0 < \alpha(\|\underline{x}\|) \leq V(\underline{x}, t)$.
- (iii) $\alpha(\|\underline{x}\|) \rightarrow \infty$ as $\|\underline{x}\| \rightarrow \infty$
- (iv) $\dot{V}(\underline{x}, t)$ is non-positive and not identically zero along any solution of the differential equation except $\underline{x} = \underline{0}$.

then the equilibrium state $\underline{x} = \underline{0}$ is UASIL and V is called a Lyapunov function.

To apply this theorem, let V be defined by:

$$V = \left[\frac{5}{3} + \cos(\theta_2) \right] \dot{\theta}_1^2 + \left[\frac{2}{3} + \cos(\theta_2) \right] \dot{\theta}_1 \dot{\theta}_2 + \frac{1}{3} \dot{\theta}_2^2 - \frac{1}{2} k_1 \theta_1^2 - \frac{1}{2} k_2 \theta_2^2 \quad (\text{B.9})$$

which is just the kinetic energy of the system from equation (B.3) plus terms which would be the potential energy of the system if the control gains were spring constants of physical springs. Hence, V corresponds to the total system energy.

Clearly $V(\underline{0}, t) = 0$, so condition (i) is satisfied. If it is required that $k_1 < 0$ and $k_2 < 0$, the only potentially negative term of V is $\left[\frac{2}{3} + \cos(\theta_2) \right] \dot{\theta}_1 \dot{\theta}_2$. The first three terms of equation (B.9) can be factored as follows:

$$\begin{aligned} & \left[\frac{5}{3} + \cos(\theta_2) \right] \dot{\theta}_1^2 + \left[\frac{2}{3} + \cos(\theta_2) \right] \dot{\theta}_1 \dot{\theta}_2 + \frac{1}{3} \dot{\theta}_2^2 = \\ & \left[\frac{5}{3} + \cos(\theta_2) \right] \left[\dot{\theta}_1 + \dot{\theta}_2 \frac{2 + 3 \cos(\theta_2) + \sqrt{9 \cos^2(\theta_2) - 16}}{10 + 6 \cos(\theta_2)} \right] \\ & \left[\dot{\theta}_1 + \dot{\theta}_2 \frac{2 + 3 \cos(\theta_2) - \sqrt{9 \cos^2(\theta_2) - 16}}{10 + 6 \cos(\theta_2)} \right] \end{aligned} \quad (\text{B.10})$$

The first factor is always positive, and since $9 \cos^2(\theta_2) - 16 < 0$, the second and third factors can be written $(x + iy)(x - iy) = x^2 + y^2 > 0$. Writing V in this way:

$$\begin{aligned} V = & \left[\frac{5}{3} + \cos(\theta_2) \right] \left[\dot{\theta}_1 + \dot{\theta}_2 \frac{2 + 3 \cos(\theta_2)}{10 + 6 \cos(\theta_2)} \right]^2 + \dot{\theta}_2^2 \frac{16 - 9 \cos^2(\theta_2)}{(10 + 6 \cos(\theta_2))^2} \\ & - \frac{1}{2} k_1 \theta_1^2 - \frac{1}{2} k_2 \theta_2^2 \end{aligned} \quad (\text{B.11})$$

By inspection of the above equation, $V(\underline{x}, t) > 0$ for all $\underline{x} \neq \underline{0}$, and V is not an explicit function of t , so V is positive definite, satisfying condition (ii) of LaSalle's Theorem. Condition (iii) is also clearly satisfied by equation (B.11).

The derivative of $V(\underline{x}, t)$ is given by:

$$\dot{V}(\underline{x}, t) = k_3 \dot{\theta}_1^2 + k_4 \dot{\theta}_2^2 \quad (\text{B.12})$$

$\dot{V}(\underline{x}, t)$ is non-positive for $k_3 \leq 0$ and $k_4 \leq 0$. If we further restrict $k_3 \neq 0$ and $k_4 \neq 0$, then the only region in which $\dot{V}(\underline{x}, t) = 0$ is the plane $\dot{\theta}_1 = \dot{\theta}_2 = 0$. Since the only trajectory which lies in this plane and satisfies the differential equation is the origin, condition (iv) of LaSalle's theorem is satisfied.

The Lyapunov function V therefore satisfies LaSalle's theorem and proves the system to be UASIL with the following constraints on the gains:

$$\begin{aligned} k_1 &< 0 \\ k_2 &< 0 \\ k_3 &< 0 \\ k_4 &< 0 \end{aligned} \quad (\text{B.13})$$

This confirms one's intuition, as it agrees with the results of linearized analysis using equations (B.6). Also, an analogy can be made between collocated PD control with negative gains and the action of mechanical springs and dampers attached to the joints. Since such components are passive they do not add energy to the system. If motion exists, the dampers remove energy, so stability follows.

B.4 Derivation of Equivalent Linear System

The next step is to introduce backlash into the system, using the input-output relationship illustrated in Figure 2.3. Note that the backlash is introduced here *between* the sensor and actuator. In a real manipulator this would correspond to measuring joint angle while controlling actuator position. In a more complicated system implementing force control, the position setpoint is determined by the forces sensed at the tip; this leads to a similar effect.

Since the goal is to detect potential limit cycles, Sinusoidal-Input Describing Functions (SIDF's) are used to define an equivalent linear system^[20]. This is an approximate method, which assumes a sinusoidal form for the system variables. The nonlinearities in the dynamical equations are replaced by linear terms, whose gains are the Fourier series coefficients for the first harmonic (higher harmonics are assumed negligible).

The accuracy of this method depends on the “filtering hypothesis,” that the open loop transfer function acts as a low-pass filter. This is often justified for equations describing physical systems. In this case the inertial nature of the plant provides filtering, but the backlash nonlinearity itself adds higher harmonics, so results will be checked by simulation.

It is assumed that both joints oscillate at frequency ω , with constant phase difference ϕ :

$$\begin{aligned}\theta_1 &= A_1 e^{i\omega t} & \theta_2 &= A_2 e^{i(\omega t + \phi)} \\ \dot{\theta}_1 &= i \omega A_1 e^{i\omega t} & \dot{\theta}_2 &= i \omega A_2 e^{i(\omega t + \phi)} \\ \ddot{\theta}_1 &= -\omega^2 A_1 e^{i\omega t} & \ddot{\theta}_2 &= -\omega^2 A_2 e^{i(\omega t + \phi)}\end{aligned}\quad (\text{B.14})$$

Any nonlinear term $f(\theta_1, \dot{\theta}_1, \ddot{\theta}_1, \theta_2, \dot{\theta}_2, \ddot{\theta}_2)$ can be replaced by $N_1(A_1, A_2, \omega, \phi) \theta_1$ or $N_2(A_1, A_2, \omega, \phi) \theta_2$ where:

$$N_1(A_1, A_2, \omega, \phi) = \frac{i}{\pi A_1} \int_0^{2\pi} f(A_1 \sin \varphi, \omega A_1 \cos \varphi, \dots, -\omega^2 A_2 \sin(\varphi + \phi)) e^{-i\varphi} d\varphi \quad (\text{B.15})$$

$$N_2(A_1, A_2, \omega, \phi) = \frac{i}{\pi A_2} \int_0^{2\pi} f(A_1 \sin \varphi, \omega A_1 \cos \varphi, \dots, -\omega^2 A_2 \sin(\varphi + \phi)) e^{-i(\varphi + \phi)} d\varphi$$

Each term in the dynamical equations is considered in turn, and an equivalent linear equation for (B.4) derived. Dividing by $e^{i\omega t}$ removes time dependence, and resolution into real and imaginary components gives the following four equations:

$$\begin{aligned}\frac{\omega k_3}{\pi} \left(\frac{b^2 - 2bA_1}{A_1^2} \right) - \frac{k_1}{2} - \frac{k_1}{\pi} \left[\sin^{-1} \left(\frac{A_1 - b}{A_1} \right) + \left(\frac{A_1 - b}{A_1^2} \right) \sqrt{2bA_1 - b^2} \right] = \\ 2\omega^2 \left[\frac{A_2 \cos \phi}{3A_1} + \frac{5}{3} + J_0(A_2) + \frac{J_1(A_2) \cos \phi}{A_1} + \cos(2\phi) J_2(A_2) \right]\end{aligned}\quad (\text{B.16})$$

$$\frac{k_1(2bA_1 - b^2)}{\pi A_1^2} - \frac{\omega k_3}{2} - \frac{\omega k_3}{\pi} \left[\sin^{-1} \left(\frac{A_1 - b}{A_1} \right) + \left(\frac{A_1 - b}{A_1^2} \right) \sqrt{2bA_1 - b^2} \right] =$$

$$2\omega^2 \left[\frac{A_2 \sin \phi}{3A_1} + \frac{J_1(A_2) \sin \phi}{A_1} + \sin(2\phi) J_2(A_2) \right] \quad (\text{B.17})$$

$$\frac{k_2 \cos \phi - \omega k_4 \sin \phi}{\pi} \left[\frac{\pi}{2} + \sin^{-1} \left(\frac{A_2 - b}{A_2} \right) + \left(\frac{A_2 - b}{A_2^2} \right) \sqrt{2bA_2 - b^2} \right] +$$

$$\frac{k_2 \sin \phi + \omega k_4 \cos \phi}{\pi} \left[\frac{2bA_2 - b^2}{A_2^2} \right] = \quad (\text{B.18})$$

$$\omega^2 \frac{A_1}{A_2} \left[-\frac{2A_2 \cos \phi}{3A_1} - \frac{2}{3} - J_0(A_2) + \frac{A_1 J_1(A_2) \cos \phi}{2} + \frac{A_1 J_3(A_2) \cos(3\phi)}{2} \right]$$

$$\frac{k_2 \sin \phi - \omega k_4 \cos \phi}{\pi} \left[\frac{\pi}{2} + \sin^{-1} \left(\frac{A_2 - b}{A_2} \right) + \left(\frac{A_2 - b}{A_2^2} \right) \sqrt{2bA_2 - b^2} \right] +$$

$$\frac{k_2 \cos \phi + \omega k_4 \sin \phi}{\pi} \left[\frac{b^2 - 2bA_2}{A_2^2} \right] = \quad (\text{B.19})$$

$$\omega^2 \frac{A_1}{A_2} \left[-\frac{2A_2 \sin \phi}{3A_1} + \frac{3A_1 \sin \phi J_1(A_2)}{2} + \sin(2\phi) J_2(A_2) + \frac{A_1 J_3(A_2) \sin(3\phi)}{2} \right]$$

where $J_i()$ are Bessel functions of the first kind of order i and b is the angular magnitude of the backlash, in radians.

B.5 Limit Cycle Prediction

When the amount of backlash b and the gains $k_1 - k_4$ are fixed, equations (B.16) through (B.19) are four nonlinear equations in four unknowns (A_1, A_2, ω, ϕ). A solution to these equations represents a limit cycle. Using ω_c and ζ to specify the gains, inspection of the equations shows that if a solution (A_1, A_2, ω, ϕ) exists for ω_c and ζ , then a solution ($A_1, A_2, C\omega, \phi$) exists for $C\omega_c$ and ζ . In other words, A_1, A_2 , and ϕ do not depend on

controller bandwidth ω_c , and ω is simply proportional to it. This allows a restatement of the solution as $(A_1, A_2, \frac{\omega}{\omega_c}, \phi)$, now only a function of backlash magnitude b and damping ratio ζ .

The equations (B.16) - (B.19) were solved numerically for a variety of damping ratios, assuming a backlash magnitude of .05 radians. A solution (limit cycle) was found in all cases, with the parameters shown in Figures B.2 - B.4:




Figure B.2: Limit Cycle Amplitude vs. Damping Ratio




Figure B.3: Limit Cycle Phase vs. Damping Ratio




Figure B.4: Limit Cycle Frequency Ratio vs. Damping Ratio

B.6 System Simulation

Control analysis using SIDF's to detect limit cycles should include simulation of the nonlinear system to test the accuracy of predictions. The SIDF method used above can save a considerable amount of simulation time by indicating where to look for limit cycles, but does not conclusively prove their existence or indicate whether they are stable or unstable.

Equations (B.4) plus the control law (B.5) and backlash were integrated numerically for two values of ζ : 0.1 and 2.0. The gain used was $k_1 = -20000$, giving controller bandwidth $\omega_c = 61.24$ rad/sec (9.75 Hz), a respectable but not far-fetched value for a real manipulator. Backlash was again set at 0.05 radians. A number of runs were performed from different initial conditions.

A single stable limit cycle developed in each case. For the $\zeta = 0.1$ case, Figure B.5 shows the joint space trajectory as the limit cycle is approached from rather distant initial conditions.




Figure B.5: Joint trajectory approaching limit cycle; $\zeta = 0.1$

Figure B.6 shows the fully developed limit cycle, after initial conditions have settled out.




Figure B.6: Joint trajectory in limit cycle; $\zeta = 0.1$

Figure B.7 shows the time variation of the joint angles in the developed limit cycle.




Figure B.7: Joint angles vs. time; $\zeta = 0.1$

Figures B.8 and B.9 show behavior in the $\zeta = 2.0$ case. Both of these figures describe the response after a considerable “settling in” period has taken place; however it is clear that the shape of each cycle is still changing. The motion seems likely to remain quasi-periodic, with stable average amplitude and frequency but variable cycle shape.




Figure B.8: Joint trajectory in limit cycle; $\zeta = 2.0$




Figure B.9: Joint angles vs. time; $\zeta = 2.0$

B.7 Accuracy of SIDF Predictions

The following table compares the results of prediction and simulation:

ζ	A1	A2	ω/ω_c	ϕ	
0.1	.1486	.3860	.7148	-.123	predicted
	.1716	.4486	.7187	-.099	simulated
	-13 %	-14 %	-5 %	-25 %	error (% of sim.)
2.0	.0280	.0332	.4496	-.259	predicted
	.0271	.0305	.3083	< 0	simulated
	+3 %	+9 %	+46 %	?	error (% of sim.)

Table B.1: Limit Cycle Parameters

The SIDF method predicted amplitudes reasonably well, was very accurate with frequency in one case but not the other, and relative phase prediction was fair in one case and indeterminate in the other (the ? in the table), due to the irregular shape of the response.

The difference in accuracy between the two cases can be explained with reference to Figures B.7 and B.9. In the former the response is clearly very close to sinusoidal, and in this case the predicted frequency was very accurate. In the latter the response is not even truly periodic, so the assumption of the filtering hypothesis is less valid.

Another test of the filtering hypothesis for this system is given by Figures B.6 and B.8. Were the motions truly sinusoidal the loci of the limit cycles would be ellipses. Figure B.6 is close, but in Figure B.8 the resemblance to an ellipse is rough.

The assumption made during equivalent linearization that the joints oscillate with the same frequency is closely supported by the simulation data; the consequent assumption that there is a constant phase relationship between them appears less valid. Since phase predictions themselves are of little interest compared to amplitude and frequency, the simplification provided by this assumption still seems justified.

Overall, the predictions were good. Using the predicted limit cycle as an initial condition, the simulation rapidly converged on the true limit cycle in each case tested. If simulation had been the only tool, much more time and effort would have been required to characterize the limit cycles.

Appendix C: Controller Listings

```

/*
                                JALoop.C
                                Joint/Actuator Controller
*/

#include <stdio.h>
#include <math.h>

#define Sf 120                    /* Stiction estimate (units = .01323 N) */
#define Cf 20                    /* Coulomb friction estimate (units = .01323 N) */
#define H .00187                /* Sample time (sec) */
#define KD .6215                /* Drive stiffness estimate (units = .00323 N/m) */
#define KDI 1.609               /* Inverse of drive stiffness */
#define DD .00279               /* Drive damping estimate (units = .0116 Ns/m) */
#define DA .06037               /* Actuator-damping estimate (units = .0116 Ns/m) */

static int x[2500],err[2500],f[2500],ua[2500];
static unsigned int t[2500];

main() {
    int i,j,k,fr,xrvel,fxd,y,xddot,vj,va,yd,u,limvl,limvh,x0,y0,f0;
    long vl;
    float xdxr,wa,wj,kt,dt,Ax,Bx,Axh,Bxh,mldcx,mlkcx,madcymakcy,xr,ft;
    FILE *tfile;

    wa = 30.;                    /* Actuator-loop bandwidth (Hz) */
    wj = 1.5;                    /* Joint-loop bandwidth (Hz) */
    printf("Enter Kt (N/m):\n");
    scanf("%f",&kt);
    printf("Enter Dt (Ns/m):\n");
    scanf("%f",&dt);
    Bx = exp(-kt*H/dt);
    Ax = 309.75*(Bx - 1.)/kt;
    Bxh = .03422*(Bx - 1.)/H;
    Axh = .03422*Ax/H;
    mldcx = .876*wj;
    mlkcx = .09419*wj*wj;
    madcymakcy = .17265*wa - DD;
    makcy = .026255*wa*wa;
    xdxr = 0;
    startup();
    printf("Enter 0 when ready:\n");
    scanf("%d",&j);
    xr = 0.;
    outpw(772,0);                /* Input reference force reading */
    do f0 = inpw(784);
        while ((f0 & 20480) != 0);
    f0 &= 4095;
    x0 = inpw(768);              /* Input reference X sensor reading */
    y0 = inpw(774);              /* Input reference Y sensor reading */
    for (i=0;i<2500;++i) {
        outpw(794,0);            /* Begin loop */
        /* Clear timer */
}

```

```

xrvel = 0; /* Calculate reference trajectory xr */
if (i<1335) xrvel = -32;
if (i<1068) xrvel = 0;
if (i<534) xrvel = 16;
xr += .0547*xrvel;
outpw(772,0); /* Input force sensor reading */
do fr = inpw(784);
    while ((fr & 20480) != 0);
fr &= 4095;
f[i] = 1.814*(f0 - fr);
xddot = xrvel + Axh*f[i] + Bxh*xdxr; /* Calculate desired x trajectory */
xdxr = Ax*f[i] + Bx*xdxr;
x[i] = 3.65*(x0 - inpw(768)); /* Input X sensor reading */
vj = inpw(776); /* Input X velocity counter reading */
if (vj != -32768) {
    if (vj < 0) vj = -(vj & 32767);
    if (vj == 0) vj = 1;
    vl = 32768/vj; /* Calculate X velocity */
    vj = vl;
    if (vl > 32767) vj = 32767;
    if (vl < -32768) vj = -32768;
    if ((vj < 2)&&(vj > -3)) vj = 0;
}
err[i] = xdxr + xr - x[i]; /* Calculate X error */
fxd = f[i] + mlcdx*(xddot - vj) + mlkcx*err[i]; /* Calculate X force */
if (fxd > 0) fxd += 10; /* Backlash compensation */
if (fxd < 0) fxd -= 10;
yd = KDI*fxd + x[i]; /* Calculate desired Y position */
y = y0 - inpw(774); /* Input Y sensor reading */
va = inpw(782); /* Input Y velocity counter reading */
if (va != -32768) {
    if (va < 0) va = -(va & 32767);
    if (va == 0) va = 1;
    vl = 136891/va; /* Calculate Y velocity */
    va = vl;
    if (vl > 32767) va = 32767;
    if (vl < -32768) va = -32768;
    if ((va < 5)&&(va > -6)) va = 0;
} /* Calculate Control */
u = DA*va + KD*(y - x[i]) + madcy*(vj - va) + makcy*(yd - y);
if (va>0) u += Cf;
if (va<0) u -= Cf;
if (va==0) {
    if (u>0) u += Sf;
    if (u<0) u -= Sf;
}
limvl = -(3278 + va); /* Impose saturation limits */
limvh = limvl + 6556;
if (u < limvl) u = limvl;
if (u > limvh) u = limvh;
u += 2048;
if (u < 0) u = 0;
if (u > 4095) u = 4095;
do t[i] = inpw(794); /* Wait for Control Time */
    while (t[i]<28);
outpw(790,u); /* Output Control Signal */
ua[i] = u - 2048;
}
outpw(790,2048);
if (j==0) { /* Write Data to File */
    printf("Writing to file test.dat\n");
    tfile = fopen("test.dat","w");
    t[0] = 0;
    ft = 0.;
}

```

```

fprintf(tfile, "Kt = %f\tDt = %f\tWa = %f\tWj = %f\n", kt, dt, wa, wj);
fprintf(tfile, "Time (sec)\tX (m)\tXerr (m)\tF (N)\tU (N)\n");
fprintf(tfile, "%f\t%f\t%f\t%f\t%f\n", t[0], x[0]/23413., err[0]/23413., f[0]*.01323,
ua[0]*.01323);
for (i=1; i<2500; ++i) {
    ft += 66.9e-6*t[i];
    fprintf(tfile, "%f\t%f\t%f\t%f\t%f\n", ft, x[i]/23413., err[i]/23413., f[i]*.01323,
ua[i]*.01323);
}
fclose(tfile);
}
printf("Done");
}

```

```

/*          Startup Routine; Allows Safe Turn-on of Power Supply          */

```

```

startup() {
    int t=1;
    printf("\nTurn on Op-Amp Power in the next 4 seconds.\n");
    outpw(794,0);
    do {t = inpw(794);
        outpw(790,2048);}
        while (t>=0);
    do {t = inpw(794);
        outpw(790,2048);}
        while (t<0);
    printf("Thank you.\n");
    return;
}

```

```

/*
                                RefLoop.C
                                Reference Design Controller
*/

#include <stdio.h>
#include <math.h>

#define Sf 120                    /* Stiction estimate (units = .01323 N) */
#define Cf 20                     /* Coulomb friction estimate (units = .01323 N) */
#define H .00187                  /* Sample time (sec) */
#define DA .06037                 /* Actuator damping estimate (units = .0116 Ns/m)*/

static int x[2500],f[2500],ua[2500];
static unsigned int t[2500];

main() {
    int i,j,k,fr,xrvel,y,yddot,vj,va,yd,u,limvl,limvh,x0,y0,f0;
    long vl;
    float wi,wo,kt,dt,Ax,Bx,madc,makc,xr,ft,ydxr,fydxr;
    float oldfydxr,tmp,z,fa,fb,fc,fd;
    FILE *tfile;

    wi = 30.;                    /* Inner-loop bandwidth (Hz) */
    printf("Enter Wo (Hz):\n");   /* Joint-loop bandwidth (Hz) */
    scanf("%f",&wo);
    printf("Enter Kt (N/m):\n");
    scanf("%f",&kt);
    printf("Enter Dt (Ns/m):\n");
    scanf("%f",&dt);
    Bx = exp(-kt*H/dt);
    Ax = 309.75*(Bx - 1.)/kt;
    z = 6.2832*wo*H/sqrt(2.);
    fa = 1. - exp(-z)*(cos(z) - sin(z));
    fb = exp(-2.*z) - exp(-z)*(cos(z) + sin(z));
    fc = -2.*exp(-z)*cos(z);
    fd = exp(-2.*z);
    madc = .17265*wi;
    makc = .026255*wi*wi;
    ydxr = 0.;
    fydxr = 0.;
    oldfydxr = 0.;
    startup();
    printf("Enter 0 when ready:\n");
    scanf("%d",&j);
    xr = 0.;
    outpw(772,0);                /* Input reference force reading */
    do f0 = inpw(784);
        while ((f0 & 20480) != 0);
    f0 &= 4095;
    x0 = inpw(768);              /* Input reference X sensor reading */
    y0 = inpw(774);             /* Input reference Y sensor reading */
    for (i=0;i<2500;+i) {
        /* Begin loop */
        outpw(794,0);          /* Clear timer */
        /* Calculate reference trajectory xr */
        xrvel = 0;
        if (i<1335) xrvel = -32;
        if (i<1068) xrvel = 0;
        if (i<534) xrvel = 16;
        xr += .0547*xrvel;
        outpw(772,0);          /* Input force sensor reading */
        do fr = inpw(784);
            while ((fr & 20480) != 0);
        fr &= 4095;
        f[i] = 1.814*(f0 - fr);
        x[i] = 3.65*(x0 - inpw(768)); /* Input X sensor reading */
    }
}

```

```

tmp = fydxr;
fydxr = fb*fydxr - fc*fydxr - fd*oldfydxr;
ydxr = Ax*f[i] + Bx*ydxr;
fydxr += fa*ydxr; /* Filter force reading */
oldfydxr = tmp; /* Calculate desired Y trajectory */
yd = fydxr + xr; /* Calculate desired Y trajectory */
yddot = xrvel + 18.27*(fydxr - oldfydxr);
y = y0 - inpw(774); /* Input Y sensor reading */
va = inpw(782); /* Input Y velocity counter reading */
if (va != -32768) {
    if (va < 0) va = -(va & 32767);
    if (va == 0) va = 1;
    vl = 136891/va; /* Calculate Y velocity */
    va = vl;
    if (vl > 32767) va = 32767;
    if (vl < -32768) va = -32768;
    if ((va < 5)&&(va > -6)) va = 0;
} /* Calculate Control */
u = DA*va + f[i] + mdc*(yddot - va) + makc*(yd - y);
if (va>0) u += Cf;
if (va<0) u -= Cf;
if (va==0) {
    if (u>0) u += Sf;
    if (u<0) u -= Sf;
}
limvl = -(3278 + va); /* Impose saturation limits */
limvh = limvl + 6556;
if (u < limvl) u = limvl;
if (u > limvh) u = limvh;
u += 2048;
if (u < 0) u = 0;
if (u > 4095) u = 4095;
do t[i] = inpw(794); /* Wait for Control Time */
    while (t[i]<28); /* Output Control Signal */
    outpw(790,u); /* Output Control Signal */
    ua[i] = u - 2048;
}
outpw(790,2048);
if (j==0) { /* Write Data to File */
    printf("Writing to file test.dat\n");
    tfile = fopen("test.dat","w");
    t[0] = 0;
    ft = 0.;
    fprintf(tfile,"Kt = %f\tDt = %f\tWi = %f\tWo = %f\n",kt,dt,wi,wo);
    fprintf(tfile,"Time (sec)\tX (m)\tF (N)\tU (N)\n");
    fprintf(tfile,"%f\t%f\t%f\t%f\n",t[0],x[0]/23413.,f[0]*.01323,ua[0]*.01323);
    for (i=1; i<2500; ++i) {
        ft += 66.9e-6*t[i];
        fprintf(tfile,"%f\t%f\t%f\t%f\n",ft,x[i]/23413.,f[i]*.01323,ua[i]*.01323);
    }
    fclose(tfile);
}
printf("Done");
}

```


References

- [1] Ahmad, S., "Analysis of Robot Drive Train Errors, their Static Effects, and their Compensations," *IEEE J. Robotics and Automation*, RA-4, no. 2, pp. 117-128, April 1988.
- [2] An, C. H., Hollerbach, J. M., "Dynamic Stability Issues in Force Control of Manipulators," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 890-896, Raleigh, NC, March 1987.
- [3] An, C. H., Hollerbach, J. M., "Kinematic Stability Issues in Force Control of Manipulators," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 897-903, Raleigh, NC, March 1987.
- [4] Appleton II, V., "Tom Swift and his Giant Robot," (New York: Grosset & Dunlap), 1954.
- [5] Asada, H., Ogawa, K., "On the Dynamic Analysis of a Manipulator and its End Effector Interacting with the Environment," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 751-756, Raleigh, NC, March 1987.
- [6] Åström, K. J., Wittenmark, B., "Computer-Controlled Systems," (New Jersey: Prentice-Hall), 1984.
- [7] Benjamin, M. H., "Design and Analysis of a Control System for the M.I.T. Precision Assembly Robot," Master's Thesis, M.I.T., January 1985.
- [8] Caine, M. E., "Chamferless Assembly of Rectangular Parts in Two and Three Dimensions," Master's Thesis, M.I.T., June 1985.
- [9] Cannon, R. H., Rosenthal, D. E., "Experiments in Control of Flexible Structures with Noncolocated Sensors and Actuators," *AIAA J. Guidance and Control*, 7, no.5 pp. 546-553, Sept.-Oct. 1984.
- [10] Cannon, R. H., Schmitz, E., "Initial Experiments on the End-Point Control of a Flexible One-Link Robot," *Int'l. J. Robotics Research*, 3, no. 3, pp. 62-75, 1984.
- [11] Canudas, C., Åström, K. J., Braun, K., "Adaptive Friction Compensation in DC-Motor Drives," *IEEE J. Robotics and Automation*, RA-3, no. 6, pp. 681-685, Dec. 1987.

- [12] DeLuca, A., Isidori, A., Nicolò, F., "Control of Robot Arm with Elastic Joints via Nonlinear Dynamic Feedback," *Proc. 24th IEEE Conf. on Decision and Control*, pp. 1671-1679, Fort Lauderdale, FL, Dec. 1985.
- [13] De Schutter, J., "A Study of Active Compliant Motion Control Methods for Rigid Manipulators Based on a Generic Scheme," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 1060-1065, Raleigh, NC, March 1987.
- [14] Desoer, C. A., Shahrz, S. M., "Stability of Dithered Non-Linear Systems with Backlash or Hysteresis," *Int'l. J. Control*, **43**, no. 4, pp. 1045-1060, 1986.
- [15] Eppinger, S. D., Seering, W. P., "Understanding Bandwidth Limitations in Robot Force Control," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 262-268, Raleigh, NC, March 1987.
- [16] Eppinger, S. D., "Modeling Robot Dynamic Performance for Endpoint Force Control," Ph.D. Thesis, M.I.T. Dept. of Mechanical Engineering, September 1988.
- [17] Flatau, D. R., "Design Outline for Mini-Arms Based on Manipulator Technology," *AI Memo No. 308*, MIT AI Laboratory, May 1973.
- [18] Foley, T. M., "NASA Considers System Changes in Effort to Cut Space Station Costs," *Aviation Week & Space Technology*, pp. 29-30, Jan. 25, 1988.
- [19] Freund, E., "Fast Nonlinear Control with Arbitrary Pole Placement for Industrial Robots and Manipulators," *Int'l. J. Robotics Research*, **1**, no. 1, pp. 65-78, 1982.
- [20] Gelb, A., VanderVelde, W. E., "Multiple-Input Describing Functions and Nonlinear System Design" (New York: McGraw-Hill), 1968.
- [21] Hogan, N., "Impedance Control: An Approach to Manipulation, Part 1: Theory, Part 2: Implementation, Part 3: Applications," *ASME J. Dyn. Systems, Meas. and Control*, 1985.
- [22] Hogan, N., "Stable Execution of Contact Tasks Using Impedance Control," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 1047-1053, Raleigh, NC, March 1987.
- [23] Hollars, M. G., "Experiments in End-Point Control of Manipulators with Elastic Drives" Ph.D. Thesis, Stanford Dept. of Aeronautics and Astronautics, May 1988.
- [24] Jacobsen, S. C., Iversen, E. K., Knutti, D. F., Johnson, R. T., Biggers, K. B., "Design of the Utah/MIT Dexterous Hand," *Proc. 1986 IEEE Int'l. Conf. on Robotics and Automation*, San Francisco, CA, April 1986.
- [25] Jilani, M. A., "Force Feedback Hydraulic Servo for Advanced Assembly Machines," M.S. Thesis, MIT Dept. of Mech. Eng., Nov. 1974.

- [26] Kazerooni, H., Houpt, P. K., Sheridan, T. B., "The Fundamental Concepts of Robust Compliant Motion for Robot Manipulators," *Proc. 1986 IEEE Int'l. Conf. on Robotics and Automation*, pp. 418-427, San Francisco, CA, April 1986.
- [27] Kazerooni, H., "Robust, Non-Linear Impedance Control for Robot Manipulators," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 741-750, Raleigh, NC, March 1987.
- [28] Khatib, O., Burdick, J., "Motion and Force Control of Robot Manipulators," *Proc. 1986 IEEE Int'l. Conf. on Robotics and Automation*, pp. 1381-1386, San Francisco, CA, April 1986.
- [29] Kokotovic, P. V., Khalil, H., O'Reilly, J., "Singular Perturbation Methods in Control: Analysis and Design," (New York: Academic Press), 1986.
- [30] Lagrange, J., "Mechanique Analytique," (Paris: Gauthier-Villars), 1788.
- [31] Loncaric, J., "Normal Forms of Stiffness and Compliance Matrices," *IEEE J. Robotics and Automation*, RA-3, no. 6, pp. 567-572, Dec. 1987.
- [32] Luh, J. Y. S., Fisher, W. D., Paul, R. P. C., "Joint Torque Control by a Direct Feedback for Industrial Robots," *IEEE Trans. Automatic Control*, AC-28, no.2, pp. 153-161, Feb. 1983.
- [33] Luh, J. Y. S., Walker, M. W., Paul, R. P., "Resolved Acceleration Control of Mechanical Manipulators," *IEEE Trans. Automatic Control*, AC-25, pp. 468-474, 1980.
- [34] Maples, J. A., Becker, J. J., "Experiments in Force Control of Robotic Manipulators," *Proc. 1986 IEEE Int'l. Conf. on Robotics and Automation*, pp. 695-702, San Francisco, CA, April 1986.
- [35] Marino, R., Nicosia, S., "On the Control of Robots with Elastic Joints," *Proc. Automatic Control Conf.*, pp. 69-70, Boston, MA, June 1985.
- [36] Mason, M. T., "Compliant Motion," *Robot Motion: Planning and Control*, (Cambridge: MIT Press), pp. 305-322, 1982.
- [37] Narendra, K. S., "Stability of Nonlinear Systems," *Nonlinear System Analysis and Synthesis: Volume 1-Fundamental Principles*, (New York: American Society of Mechanical Engineers), 1978.
- [38] Ogata, K., "Modern Control Engineering," (New Jersey: Prentice-Hall), 1970.
- [39] Paul, R. P. C., "Problems and Research Issues Associated with the Hybrid Control of Force and Displacement," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 1966-1971, Raleigh, NC, March 1987.
- [40] Raibert, M. H., Craig, J. J., "Hybrid Position/Force Control of Manipulators," *ASME J. Dyn. Systems, Meas. and Control*, 102, pp. 126-133, June 1981.

- [41] Roberts, R. K., Paul, R. P., Hillberry, B. M., "The Effect of Wrist Force Sensor Stiffness on the Control of Robot Manipulators," *Proc. 1985 IEEE Int'l. Conf. on Robotics and Automation*, pp. 269-274, St. Louis, MO, March 1985.
- [42] Salisbury, J. K., "Active Stiffness Control of a Manipulator in Cartesian Coordinates," *Proc. 19th IEEE Conf. on Decision and Control*, Dec. 1980.
- [43] Schultz, D. G., Melsa, J. L., "State Functions and Linear Control Systems," (New York: McGraw-Hill), 1967.
- [44] Slotine, J.-J. E., Hong, S., "Two-Time Scale Sliding Control of Manipulators with Flexible Joints," *Proc. 1986 American Control Conf.*, Seattle, WA, 1986.
- [45] Spong, M. W., Thorp, J. S., Kleinwaks, J. M., "The Control of Robot Manipulators with Bounded Input," *IEEE Trans. on Automatic Control*, AC-31, no. 6, June 1986.
- [46] Spong, M. W., Khorasani, K., Kokotovic, P. V., "An Integral Manifold Approach to the Feedback Control of Flexible Joint Robots," *IEEE J. Robotics and Automation*, RA-3, no. 4, pp. 291-300, Aug. 1987.
- [47] Stepien, T. M., Sweet, L. M., Good, M. C., Tomizuka, M., "Control of Tool/Workpiece Contact Force with Application to Robotic Deburring," *IEEE J. Robotics and Automation*, RA-3, no. 1, pp. 7-18, Feb. 1987.
- [48] Sweet, L. M., Good, M. C., "Redefinition of the Robot Motion-Control Problem," *IEEE Control Systems Magazine*, pp. 18-25, Aug. 1985.
- [49] Takegaki, M., Arimoto, S., "A New Feedback Method for Dynamic Control of Manipulators," *ASME J. Dyn. Systems, Meas. and Control*, 103, pp. 119-125, June 1981.
- [50] Townsend, W. T., Salisbury, J. K., "The Effect of Coulomb Friction and Stiction on Force Control," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 883-889, Raleigh, NC, March 1987.
- [51] Townsend, W. T., "The Effect of Transmission Design on Force-Controlled Manipulator Performance," Ph.D. Thesis, M.I.T. Dept. of Mechanical Engineering, April 1988.
- [52] Whitney, D. E., "Force Feedback Control of Manipulator Fine Motions," *ASME J. Dyn. Systems, Meas. and Control*, pp. 91-97, June 1977.
- [53] Whitney, D. E., "Quasi-Static Assembly of Compliantly Supported Rigid Parts," *Robot Motion: Planning and Control*, (Cambridge: MIT Press), pp. 439-471, 1982.
- [54] Whitney, D. E., "Historical Perspective and State of the Art in Robot Force Control," *Proc. 1985 IEEE Int'l. Conf. on Robotics and Automation*, pp. 262-268, St. Louis, MO, March 1985.

- [55] Whitney, D. E., "Real Robots Don't Need Jigs," *Proc. 1986 IEEE Int'l. Conf. on Robotics and Automation*, pp. 746-752, San Francisco, CA, April 1986.
- [56] Wlassich, J. J., "Nonlinear Force Feedback Impedance Control," S.M. Thesis, MIT Dept. of Mech. Eng., 1986.
- [57] Youcef-Toumi, K., Li, D., "Force Control of Direct-Drive Manipulators for Surface Following," *Proc. 1987 IEEE Int'l. Conf. on Robotics and Automation*, pp. 2055-2060, Raleigh, NC, March 1987.
- [58] "Brushless DC Servo Systems Catalog," *General Numeric*, Elk Grove Village, IL, 1987.
- [59] "Maxon DC Motor Catalog," *Maxon Precision Motors, Inc.*, Burlingame, CA, 1987.