

Characterization and Control of an Interactive Robot

by

Stephen Paul Buerger

B.M.E., University of Dayton (1999)


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

Master of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2001



BARKER

© Massachusetts Institute of Technology 2001. All rights reserved.

Author
Department of Mechanical Engineering
January 9, 2001

Certified by
Neville Hogan
Professor
Thesis Supervisor

Accepted by
Ain A. Sonin
Chairman, Department Committee on Graduate Students

Characterization and Control of an Interactive Robot

by

Stephen Paul Buerger

Submitted to the Department of Mechanical Engineering
on January 9, 2001, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

The field of rehabilitation robotics demands robots that can interact softly with the patients they treat while also bearing substantial loads. This thesis discusses the characterization of a screw-driven module that, when added to an existing therapy robot, will expand the workspace of this robot from two dimensions to three, broadening the effectiveness of the therapy that can be delivered. The results of the characterization are presented and provide a thorough understanding of the control electronics, the conversion of control from the signal domain to mechanical power, and the properties of the mechanical transmission. A detailed study of friction is included, as friction was found to be a dominant factor in design performance. The use of active control methods to compensate for friction and to improve performance is considered and discussed. Suggestions are presented concerning the further examination of these types of controllers to facilitate the design of robots for rehabilitation and interactive machines in general.

Thesis Supervisor: Neville Hogan
Title: Professor

Acknowledgments

The work described here, though enjoyable in its own right, has been made infinitely more so thanks to my contact with many great people. Foremost among these is my advisor, Neville Hogan. Neville's prominence in his field is no accident, and his insight is rivaled only by his enthusiasm. A few minutes with Neville has never failed to revitalize me and send me in the right direction.

Brandon Rohrer has been indefatigable in his willingness and ability to answer my many, many questions. In fact, he has refined this art to the point that all I must do is sit down to write him an email and the answer often descends upon me. Like Brandon, Max Berniker helped me immensely with this thesis and has greatly extended my understanding of our field through our many discussions, all this while captaining our immensely successful hockey team and sparking a hot new recording group.

Special thanks to Kris Jugenheimer and Dustin Williams, with whom I designed the hardware discussed here. Without their clever design work, the characterization could never have happened. Igo Krebs provided useful guidance thanks to his experience in getting these robots successfully into the hospitals. Lori Humphrey can solve almost any problem with a smile. Jerry Palazzolo, James Celestino, Sue Fasoli, and Phil Tang round out a top notch research group.

Thanks also to all the members of the Hunter lab for the use of all their resources and especially for all the Newman Lab camaraderie. They have grown too numerous to list completely, but Bryan Crane and Patrick Anquetil have become close friends and confidants. Bryan and I have faced many of the same learning curves and adjustments, and having somebody in the same boat is always appreciated. Patrick can always pick me up when I am down. Keng "KH" Lim and Luke Sosnowski are sorely missed. Thanks also to John, Peter, James, Robert, and all the rest.

My parents have been very supportive and are always there to talk, even when I know they just want me to come home. Thanks to all my family and my tremendous set of friends from UD, MIT, and everywhere else, especially to Liz for all her support. I miss them all greatly and am constantly hoping they will continue to visit me here on the East Coast.

This work has been partially supported by a fellowship from the Department of Defense, and by the Winifred Masterson Burke Medical Research Institute.

Contents

1	Introduction	17
1.1	MIT-MANUS and MIT-MANUS II	19
1.2	A 3 Degree of Freedom Therapy Robot	21
2	Design Process	25
2.1	Major Requirements	25
2.1.1	Workspace Dimensions and Endpoint Forces	25
2.1.2	Minimum Endpoint Impedance	27
2.1.3	Other Important Requirements	30
2.2	Configuration Selection	31
2.2.1	New 3 Degree of Freedom Designs	31
2.2.2	Advantages of MANUS	31
2.2.3	SCARA-based designs for 3 degrees of freedom	32
2.2.4	Adaptability of MANUS	32
2.2.5	Transmission Selection	33
2.3	Screw-Driven Design	34
3	System Operation	37
3.1	Hardware Overview	37
3.1.1	Rollnut	37
3.1.2	Linear Bearing	40
3.1.3	Mounting Support	42
3.1.4	Flexible Coupling	43

3.1.5	Alignment Difficulties	43
3.1.6	Servomotor	47
3.1.7	Servoamplifier	48
3.1.8	Position Sensor	48
3.1.9	Force Sensor	48
3.1.10	Adjustable Handle Module	49
3.1.11	Vertical Locking Device	49
3.2	Software Overview	49
3.3	Typical System Operation	50
3.4	Studying the system	53
4	Signal to Mechanical Power Subsystem	55
4.1	Amplifier and Actuator analysis	55
4.1.1	Low Frequency Static Testing	55
4.1.2	Amplifier Frequency Response	57
4.1.3	Torque frequency response	62
4.1.4	Limitations of delay information	67
4.2	Subsystem conclusions	69
5	Transmission Subsystem	71
5.1	Parameters to be identified	71
5.1.1	The importance of friction	74
5.1.2	Why do we need to know?	75
5.2	Friction Identification	76
5.2.1	Understanding Machine Friction	76
5.2.2	Model Simplifications	85
5.2.3	Experiment Design	87
5.2.4	Experimental Results	90
5.2.5	Model for Friction in the Screw Module	109
5.3	Inertia Identification	116
5.3.1	Estimation	118

5.3.2	Experiment	119
5.3.3	Experimental Results	121
5.4	Subsystem Conclusions	124
6	Controls Analysis	125
6.1	Linear Impedance Controller	125
6.2	Friction Compensation Techniques	131
6.2.1	Feedforward Controllers	131
6.2.2	Feedback Controllers	133
6.3	Applying Friction Compensation to an Impedance Controller	134
6.3.1	Impedance and Interaction	134
6.3.2	Feedback Control	135
6.4	Compensation Applied to the Module	136
6.4.1	Negative Software Damping	136
6.4.2	Feedforward Friction Compensator	138
6.4.3	Performance Enhancement with Force Feedback	140
6.4.4	Controller for the Screw Module	147
7	Conclusions	149
7.1	A Framework for Further Investigation	150
	Appendix A Compliant Mounting Analysis	153
A.1	Model for Compliant Mounting	153
A.1.1	Driving Frequencies	154
A.2	Case 1: Motor and mount piece move with screw	155
A.3	Case 2: Motor and mount piece grounded	156
A.3.1	Damping	157
A.4	Case 3: Tuned Vibration Absorber	157
A.5	Conclusions	161
	Appendix B Software Timing Study	163
B.1	Testing	163

B.2 Intentional Delays	164
Appendix C Models for Frequency Response	167
C.1 Simulated system	167
C.2 Amplifier Characterization Model	167
C.3 Actuator Characterization Model	169
Appendix D Conference Paper Detailing the Design	171
Bibliography	179

List of Figures

1-1	MIT-MANUS II therapy robot.	20
1-2	Patient working with MIT-MANUS	20
2-1	Solid model of the new module for vertical motion.	34
3-1	Screw/nut interface of a Rollnut.	39
3-2	Rollnut assembly.	40
3-3	Reaction force and torque at nut	41
3-4	Linear ball bearing, rail and carriage.	42
3-5	Structural support/mount piece.	43
3-6	Helical beam coupling joining motor shaft to screw.	44
3-7	Sources of screw binding	45
3-8	Grommet used to provide compliance in top bearing mounting.	46
3-9	System schematic.	51
3-10	Virtual spring and damper	52
3-11	System block diagram	54
4-1	Locked rotor testing	56
4-2	Torque vs. commanded current	57
4-3	Commanded current vs. desired torque, for linearization	58
4-4	Torque vs. current, linearized	58
4-5	Amplifier frequency response, DSA	60
4-6	Frequency Response data as in Figure 4-5 with fitted model.	61
4-7	Amplifier frequency response, sampled at 2 kHz	62

4-8	Amplifier frequency response, data and model	63
4-9	Actuator frequency response	64
4-10	Actuator frequency response, data and model	65
4-11	Phase vs. frequency, linear scale	66
4-12	Actuator frequency response, sampled at 1 kHz	67
4-13	Ideal sampling of a unity gain plant.	68
4-14	Sampling a system with a partial sample delay.	69
5-1	Translational free body diagram	73
5-2	Asperities in contact	78
5-3	Anticipated friction velocity dependence	79
5-4	Experimental setup, constant velocity testing	88
5-5	Interface for force transmission for friction testing against gravity. . .	89
5-6	Double ball-joint interface	89
5-7	Friction measurements in the frequency domain	92
5-8	Friction measurements in the frequency domain, closeup	92
5-9	Optimal low pass filter	94
5-10	Velocity dependence of friction at bottom of screw	95
5-11	Velocity dependence of friction at top of screw	96
5-12	Position dependence of friction	98
5-13	Position dependence of friction, range of speeds	99
5-14	Friction force vs. position and velocity	100
5-15	Repeatability of friction measurements	101
5-16	Velocity dependence of friction, downward	102
5-17	Velocity dependence of friction, downward	103
5-18	Position dependence of friction, downward	104
5-19	Friction in both directions	104
5-20	Force with and against gravity, positive direction	106
5-21	Force with and against gravity, negative direction	106
5-22	Friction force with gravity	107

5-23	Friction force vs. position for gravity identification	109
5-24	Gross fit of position dependence, upward	112
5-25	Fit for position dependence of friction, upward	112
5-26	Fit for velocity dependence of friction	114
5-27	Friction minus friction model, upward	114
5-28	Fit for position dependence of friction, downward	115
5-29	Fit for velocity dependence of friction, downward	117
5-30	Friction minus model	117
5-31	Low pass filter	120
5-32	Acceleration and mass times acceleration curves	122
5-33	Acceleration and mass times acceleration curves, closeup	122
5-34	Calculated inertia vs. time	124
6-1	Force vs. position, PD controller, 100 N/m	127
6-2	Force vs. position, PD controller, 1000 N/m	128
6-3	System under PD control interacting with human arm	130
6-4	Model appearance for desired system behavior	130
6-5	Negative viscous damping	137
6-6	System under PD control with friction and gravity compensators. . .	138
6-7	Block diagram, PD controller and friction estimator	139
6-8	Performance with friction compensator	140
6-9	Stiction that remains with friction compensator	141
6-10	Performance with force feedback controller	143
6-11	Simulation for force feedback	144
6-12	Simulated results using PD controller	145
6-13	Simulated results using force feedback	145
A-1	Model for system with compliance	154
A-2	Moment of inertia about pivot vs. nut location, Case 1.	155
A-3	Natural frequency vs. nut location, Case 1.	156
A-4	Natural frequency vs. nut location, Case 2.	157

A-5	Model for system with tuned vibration absorber.	158
A-6	Bond graph of tuned vibration absorber	159
A-7	Frequency response with tuned vibration absorber	160
C-1	Block diagram of simulation models	168

List of Tables

3.1	Approximate timing of major software operations	50
5.1	Estimated contributions to endpoint mass, by component.	118
B.1	Approximate timing of software operations in each sampling period. .	166

Chapter 1

Introduction

Recovery from brain injury is a major concern of medical researchers. Stroke is a leading cause of disability in the United States [23] and a major health problem worldwide. With an aging population, particularly in North America, this problem is not likely to subside.

Many patients suffer as a consequence of brain injury partial or complete loss of motor control in their limbs. Such paralysis is manifested in several different ways but often appears as hemiplegia, a paralysis concentrated on one side of the body. Some patients eventually recover partial or complete use of the affected limbs; the recovery process is typically on the order of weeks and months, and can vary significantly from patient to patient.

Recent research has shown that recovery can be enhanced through the administration of physical therapy. Furthermore it has been shown that the use of robotic therapy aids can speed the recovery process. Aisen et al. [2] studied recovering stroke patients and showed that the performance of an experimental group, that received robot therapy in addition to the typical therapy regimen, improved by a greater amount than the performance of a control group in a statistically significant sense.

The importance of the results from Aisen et al. must not be overlooked. These findings are significant as they demonstrate the importance of therapy and the viability of robotic therapy. Even beyond these important realizations, the work potentially points toward more revolutionary ideas in motor recovery. Clearly therapy helps to

speed and complete recovery. However most therapy techniques in use in today's hospitals and clinics are not founded in the science of human motion and motor recovery, but are largely derived from intuition about the types of movements that are useful, or are specifically targeted to the completion of useful tasks for the patient. One might hypothesize that different types of therapy might affect recovery in different ways; thus one might discover techniques which are preferred for promoting recovery. As current understanding of the human motor control system is quite limited, the process of motor recovery could both assist in and benefit from the development of a better understanding of the biological processes at work. Such a full understanding could potentially lead to the creation of specific therapy routines providing for optimal recovery and targeting very specific types of lesions and patients.

The process proposed in the preceding paragraph is yet in its infancy, and is presented largely to provide the reader with a more fully developed "big picture" understanding of the importance and significance of the research described herein. The apparent potential of these ideas is great; as more information about human motor recovery is gathered, these ideas can be evaluated in the context of a greater understanding.

The uses of robot therapy are certainly not restricted to the long term vision of targeting therapy for optimal recovery. The technology has very immediate uses in complementing the work of physical and occupational therapists. Robotic therapy units are conceived of not to replace therapists but to expand their capabilities. With a robotic therapist, a patient may one day be able to receive therapy between sessions with his human therapist; this could be useful as patients tend to tire quickly but could potentially receive therapy several times a day. Perhaps a robot unit could be dispatched to the home of a patient recovering from brain injury so he could receive portions of his therapy without staying in or traveling to a hospital.

Robot therapy can also be extremely useful as a measurement tool. Clinical measures of patient performance are largely based on subjective measures performed by the human therapist. The therapist tests the patient's strength and coordination in attempting one of several tasks. These measures, while widely accepted in the field,

are largely unscientific and are heavily susceptible to human error and bias. Conversely, the robot can be used to quantitatively measure both the motion and the forces that the patient is producing to a high degree of accuracy. This data can be stored and analyzed to provide detailed information about patient recovery. Developing consistent and successful ways of processing this data to provide information about the patient is an important part of robot therapy research.

The following section will describe the robot used in the clinical trials for robotic therapy.

1.1 MIT-MANUS and MIT-MANUS II

MIT-MANUS and MIT-MANUS II are the alpha and beta prototypes, respectively, of a first generation of robots designed to administer therapy to patients recovering from brain injury. The two robots differ only in subtle points of design.

The MANUS design [7, 1] utilizes a 5-bar Selective Compliance Assembly Robot Arm (SCARA) configuration to deliver power from its two stationary brushless DC servomotors to the point of patient interaction, a handle located in a planar workspace. MANUS II is shown in Figure 1-1.

When a patient works with the robot, his hand and forearm are strapped into a cradle attached to the robot's handle. The patient is seated at a table facing the robot and a computer monitor. The monitor provides the patient with visual feedback of his performance and goals. The patient plays video games, generally attempting to move his hand from point to point in a planar workspace parallel to the tabletop, as dictated by the game. The MANUS robots help in rehabilitation of the shoulder and elbow muscles but only for a certain set of movements in a horizontal plane. Figure 1-2 shows a patient playing a game on the robot.

In the period immediately following a stroke, a patient is often unable to move his impaired limb at all. In this early stage of recovery, the robot moves the patient's hand from point to point as the patient tries to complete the same moves. As the patient's condition improves, so does his performance, and the interaction with the




Figure 1-1: MIT-MANUS II therapy robot.




Figure 1-2: A patient receiving therapy by playing an interactive video game on MIT-MANUS.

robot becomes more evenly divided; he is moving the robot as much as the robot is moving him. Late in recovery, the robot is often used without power as the patient pushes it around from target to target. The robot can also be used to provide resistance once the patient can move on his own, to help with strength training.

To deliver therapy in the way described in the preceding paragraph, a robot with a soft, compliant feel is required. In the MANUS design, this is achieved through the use of backdrive-able hardware and impedance control. The SCARA configuration, used with direct drive motors, is inherently backdrive-able. The importance of such a low-impedance interaction is addressed extensively throughout this thesis. Impedance control [19] provides a soft, stable environment for interaction. Using one implementation of impedance control the environment feels an apparent spring and damper between the endpoint position and the desired endpoint position at any given time.

The MANUS design has provided a platform for the successful robotic therapy research described above.

1.2 A 3 Degree of Freedom Therapy Robot

As mentioned above and in the work by Aisen et al., the administration of robotic therapy produces significantly greater recovery in patients than that seen in those who do not receive robot therapy. It is important to note, however, that the improvements are limited to the specific muscle groups exercised by the therapy; there is no significant generalization of improvement to other muscle groups [2]. Thus in order to provide therapy that will lead to full recovery, it is necessary to develop additional robot modules to target other specific muscle groups.

The first and second incarnations of MIT-MANUS, as described in the preceding section, provide for motion exercising the shoulder and elbow muscle groups through movements in a planar workspace oriented parallel to a tabletop. This workspace provides an excellent framework to train and study a controlled set of movements to evaluate the effectiveness of therapy and robotic therapy. The movements trained,

however, certainly do not encompass a full set of the useful motions for a normally functioning human being. Even for the elbow and shoulder groups, more extensive training is useful to learn tasks such as feeding and reaching, tasks that are essential for independence. It is critical that the patient eventually recover the ability to move against gravity, and to carry the weight of his arm in addition to other objects.

With this need in mind, a new module was designed to augment the MANUS configuration by providing for motion in a third degree of freedom. The planar workspace was expanded vertically into a three-dimensional workspace. No significant loss of planar function was suffered, and the capability to move up and down without the planar confines of MANUS was provided. The module offers greatly expanded therapy possibilities and will speed shoulder and elbow recovery.

A device was created that is mechanically simple and that meets the desired goals. The module was tested and analyzed extensively. The major limitations to module performance were identified and steps were taken to compensate for these limitations.

This thesis discusses the development and implementation of a prototype of the new module design. The details of the design process are discussed elsewhere [22](see Appendix D); this document primarily focuses on the characterization, calibration, and identification of the robotic system, and of measures taken to bring the new module's performance into the desired specification range.

- Chapter 2 provides a brief discussion of the design and of a few of the major design decisions and the factors that motivated them. Also included is discussion of modifications made to the design because of revelations in the course of fabrication.
- Chapter 3 gives an overview of the major hardware and software components, and provides a detailed discussion of how the system works.
- Chapter 4 discusses the characterization and identification of the portion of the system providing transmission from the theoretical controller command to mechanical power at the motor shaft.

- Chapter 5 describes the process and results of examining the mechanical transmission in the device, from the motor to the port of contact with the patient. This includes an overview of the literature concerning machine friction, including theory and analysis, as this topic is central to study of the module.
- Chapter 6 discusses several approaches to impedance control of the module, including compensation techniques for the improvement of system performance through active control rather than mechanical redesign.
- Chapter 7 provides conclusions drawn from the work including suggested improvements for this specific design and recommendations on how to build on this work.

Chapter 2

Design Process

This chapter discusses in summary the design of a therapy robot for training movement in a three-dimensional workspace. The design process for this device is described in greater detail in Appendix D [22]; this chapter simply attempts to provide the background necessary to understand the design and the motivation for the work described here. Fundamental requirements for the device were defined. Based on these requirements a large number of concepts for the new design were generated. These designs were considered and evaluated on the basis of their fulfillment of the requirements. The process of selecting a suitable configuration was an iterative one involving brainstorming, basic calculations, and systematic idea processing. This chapter is intended not to deductively justify the choice of design but to describe the general evolution of the design from a set of requirements and many unrelated ideas to the one idea actually developed.

2.1 Major Requirements

2.1.1 Workspace Dimensions and Endpoint Forces

The robot discussed here is, from conception, a three-dimensional analogue to the two-dimensional MANUS design. As such, the new robot is expected to approach or exceed the performance of MANUS in the planar workspace, and to provide analogous

performance in the spatial area above the workspace.

MANUS nominally offers a planar workspace that measures 15 by 18 inches (0.381 by 0.4572 m). The new device was targeted to offer a comparable range of motion vertically, nominally 18 inches (0.4572 m).

In order to offer sufficient stiffness, MANUS was designed to provide sufficient forces to counteract the forces that a typical patient can produce in each direction. Charnnarong provides a discussion of the determination of the targeted endpoint force for the MANUS-I robot [7]. He points out that most patients are weaker than an average weak person as documented in [11]. Charnnarong decides on an approach of designing for slightly greater power than necessary, and limits endpoint forces to approximately 10 lbf (44.5 N). This value is sufficient for most therapy, but does not provide enough power to overcome extreme patient conditions such as spasticity and tone. 10 lbf (44.5 N) is enough to offer performance comparable to MANUS and is therefore an acceptable minimum, but greater force capabilities may prove advantageous.

The new robot is required to meet similar force requirements in the plane. Furthermore the robot must offer significant force in the vertical direction. This force must be enough to compensate for the weight of a patient's arm as well as to offer sufficient forces for effective therapy. Vertical force producing capabilities are comparable to forces in the plane [11]. Following Charnnarong's reasoning, the robot should be able to provide at least 10 lbf (44.5 N) in addition to supporting the patient's arm weight.

Effective arm weight is a fairly complicated function not only of arm dimensions and mass, but also of arm configuration. Because of the large number of variables, some of which are difficult to estimate, a few simple experimental measurements were used to empirically determine a reasonable estimate of arm weight. Measurements were made both on healthy people attempting to minimize muscle action and on a recovering patient with a flaccid arm, in a variety of arm configurations. The patient was obese, so the measurements give an idea of the "worst case".

To measure effective arm mass, the patient's hand was placed in a brace, and

suspended from a fixed tripod via a spring scale. A force was then read from the spring scale. All measurements indicated a mass less than 2 kg, leading to a force less than 20 N.

Based on arm mass and desired force estimates, the target maximum endpoint force was set at approximately 65 N upward and 45 N downward.

2.1.2 Minimum Endpoint Impedance

The need for low impedance

The importance of providing the possibility of low endpoint impedance for this type of robot is worth emphasizing before specifically addressing the characteristics of any particular hardware. This point has great significance in distinguishing the technology discussed here from typical modern robots. Standard industrial robots are restricted to extremely high endpoint impedances. Often there is so much friction in the heavily geared transmission of such a robot that the manipulator cannot be moved without its actuators commanding its position. For a normal person this is an extremely “hard” feel, with virtually no compliance. Such robots are generally position or trajectory controlled. If an object obstructs their commanded path, they push it out of the way; if the object is “stronger” than the robot, then the robot breaks (or more likely triggers some sort of self-shutdown).

This feel is far from that desired for a therapy robot. A therapy robot should provide a soft environment with which the patient can interact. The robot should gently encourage the patient to move in the desired path, but should not generally interfere with the patient’s own efforts to move. There is significant evidence to suggest that human upper-limb movement is organized based on hand motion in body-centered coordinates, rather than in jointspace [14]. If a key portion of the process of human motor control is achieving and potentially observing movement in an external coordinate frame, it is important not to infringe on the ability to move when encouraging motor recovery. It would be counter-productive to interfere with a patient’s efforts to move, and doing this would likely result in ineffective therapy.

Ideally a therapy robot should offer the possibility of an endpoint impedance ranging from nominally zero to stiff enough that even a strong patient has difficulty moving it. This scenario offers the flexibility of administering therapy and performing evaluation in a variety of virtual environments, at the discretion of the therapist and attending physician. This goal permeates the entire development process of a therapy robot and is kept in mind at all stages of the design, characterization, and testing.

Defining Impedance

With the possibility of small endpoint impedance identified as a critical design goal, it becomes important to specifically define what this means. Minimizing endpoint impedance means, in the context of the design of this robot, minimizing both endpoint friction and effective endpoint mass.

For a port of interaction in a mechanical system, impedance can be understood as a relationship between the force and motion, specifically velocity, at the port. A port that acts as an impedance provides some force in response to a motion imposed on it. Such a port must interact with an admittance, a port that allows motion in response to some force. In our case, the environment or patient is viewed as an admittance, and the robot provides a certain endpoint impedance [19].

Impedance is dictated by the effective mechanical properties of a system at the port of interaction. For the linear case, this can be shown by considering, for example, a simple second order system described by the following equation:

$$M\ddot{x} + b\dot{x} + kx = u \tag{2.1}$$

where u is input force. In the Laplace domain, with the Laplace variable s , this system's transfer function form can be represented as:

$$\frac{U}{X} = Ms^2 + bs + k \tag{2.2}$$

Or, substituting $V = sX$ where V is velocity,

$$\frac{U}{V} = Ms + b + \frac{k}{s} \quad (2.3)$$

As impedance is defined as the left half of equation 2.3, the force necessary to produce a desired motion, it can clearly be seen that for the linear case any increase in M or b will result in a higher impedance. It is also important to minimize k , the stiffness about an operating point, but this is generally controlled in software for the class of designs that has no strongly preferred orientation. This is the case for the class of designs discussed here.

Although equation 2.3 represents a linear system with b a constant damping coefficient and M a constant mass, note that this idea can be extended to nonlinear friction and state-dependent mass as well. This case effectively means that b and/or M is a function of time or the state. Clearly it can be seen that for minimum impedance, it is desirable to minimize b and M regardless of time or state.

Minimum Friction

In conjunction with the minimization of endpoint mass, discussed in the next section, friction minimization is essential to the new robot. If one seeks only to minimize friction, it is sometimes possible to replace a large amount of the potential friction with larger inertia, for example by eliminating geartrains in favor of large direct drive actuators. In a design where both friction and inertia must be minimized, however, this tradeoff must be tempered and carefully considered.

An overview of machine friction is reserved for Chapter 5, where it is presented in the context of the chosen design. Here it must only be noted that the minimization of friction was of primary concern from the beginning of the design process.

Minimum Mass

As the patient is intended to manipulate the endpoint of the robot by pushing and pulling it, minimizing the effective endpoint mass is important in keeping impedance

small. This can clearly be seen from equation 2.3 or from Newton’s second law of motion, which can be stated as:

$$\sum F = M\ddot{x} \quad (2.4)$$

Clearly it takes less force to accelerate a smaller mass. As stroke patients are not expected to produce substantial forces in the early stages of recovery, it is important to keep the mass against which they must work to a minimum.

2.1.3 Other Important Requirements

A number of requirements beyond the endpoint force requirements and the minimization of endpoint impedance were carefully considered in the design of the new robot. Several of the most important are highlighted here.

Safety is always of primary concern, especially in a clinical environment such as that in which this robot must operate. This robot must be equipped with multiple safety systems to ensure safe operation and stable shutdown should the need arise. The design must be free of pinch points and must guarantee stability, even in the case of operator error.

The interface to the patient should have a “smooth” feel. The patient should feel comfortable in dealing with the robot, and should not get the impression of dealing with a mechanical device. To this end, it is important to keep any backlash and play in the system to a minimum, and to ensure that the system can actuate smoothly at extremely slow speeds.

It is preferable to create a system that utilizes electromechanical actuators for portability and easy use in a clinical environment. This requirement could be circumvented but only with great difficulty.

MANUS-II is a compact, modular design that allows for easy transportation. It also keeps much of its bulky hardware away from the patient, so as not to appear too intimidating. It is preferable to preserve these qualities in the new robot.

2.2 Configuration Selection

2.2.1 New 3 Degree of Freedom Designs

Several configurations were proposed that involved a substantial departure from the style of MANUS and MANUS-II. These concepts used many different types of actuators and transmissions to produce motion in three dimensional space.

Each of the new configurations considered seemed extremely complex, and offered few advantages over many of the other designs. Complexity and cost, without any obvious advantages, were the main factors in eliminating such designs.

2.2.2 Advantages of MANUS

While the SCARA design used in MANUS is limited to two dimensional motion, it offers many properties advantageous to the application. The 5-bar closed chain mechanism permits motion over a fairly large area with stationary actuators. Furthermore, within certain bounds of this area, the effect of the inherent nonlinearities in inertia and damping are not significant; this property permits the use of the specified workspace, much smaller than the total range of manipulator motion, with negligible deviations in inertia and damping. Foster's [15] results particularly illustrate this point, showing less than a 70 percent variation in endpoint inertia with direction. Coulomb friction and viscous damping variation is a bit more substantial, but these values are so low that it causes few real problems.

The long arms of the linkage require precise sensor resolution, but also reduce the effective endpoint mass substantially by providing mechanical advantage in backdriving the actuators. The result is a large lever arm to aid in overcoming bearing and actuator friction as well as actuator inertia.

A final advantage of the SCARA design is suggested by its name: Selective Compliance Assembly Robot Arm. While the device is extremely compliant (has very low impedance) in the plane of its actuation, it is extremely rigid in the third dimension. Thus the device has no problems bearing any vertical forces the patient places on it.

2.2.3 SCARA-based designs for 3 degrees of freedom

The several advantages of the SCARA configuration described above proved difficult to ignore. Many of these advantages are desired in the new device as well. There is no clear 3 dimensional analogue to a SCARA configuration; in other words there is no closed chain mechanism known to the author that provides the desired three degrees of freedom and the same advantages of a SCARA. Thus a class of design concepts developed that included the idea of employing a SCARA for two degrees of freedom and an additional actuation technique for the third degree of freedom.

One subset of such designs included the idea of moving the base of the SCARA to provide the third degree of freedom. Ideas included moving the entire linkage and its actuators vertically, or rotating the entire structure. Each of these ideas involved the need to move extremely massive actuators, greatly increasing endpoint inertia and requiring an extremely powerful third actuator. An alternative idea considered moving only the SCARA links vertically on shafts extending from the actuators, while keeping the actuators stationary. This idea was eliminated largely because of concerns about binding due to the large torques that would be transmitted to the sliding parts, and because the inertia was still substantial.

Another subset of SCARA-based designs included designs in which a new module is mounted at the handle of a robot resembling the existing robot, and is carried about the plane. This module provides the vertical motion. These designs were attractive because of the idea of modularizing the robot, of potentially meeting the desired criteria by simply adding a piece of hardware to the planar design.

2.2.4 Adaptability of MANUS

To consider the idea of adding a module at the endpoint of MANUS, it was necessary to evaluate the capability of the existing robot to support such a new device, and to bear the expected loads associated with vertical actuation.

Charnnarong's [7] work provides the calculations used in the structural design of the MANUS SCARA. Using an expected maximum vertical endpoint force of 20 lb

(89 N), he applies a design factor of 2.5. Then in designing the links to handle the maximum bending moment, he chooses a material that provides an additional design factor of 2.825, giving a total design factor of just over 7. The joints and bearings are designed with similar loads and design factors in mind.

Preliminary calculations suggested that an endpoint-mounted module might weigh approximately 12 lb (5.5 kg). The force from this weight, when added to the expected endpoint force of 20 lb (89 N), suggests that a maximum force of 32 lb (142 N) would be expected. This indicates that MANUS, or a robot like it, would be able to bear the added load while maintaining a design factor of 4.4. Thanks to the conservative design of the original device, it is reasonable to suppose that such a design could be implemented on one of the MANUS robots with minimal change or redesign.

2.2.5 Transmission Selection

Several different ideas for the endpoint-mounted module were considered. As stated above, the module must provide vertical motion while minimizing mass so as to be carried around in the plane with minimal effort. The effective endpoint mass in the plane, documented experimentally by Foster [15], would increase by the weight of the entire module.

A linear motor was considered as a direct drive option for the module at the end. The limiting factor for a linear motor in this design is the mass of the module. In order to get a motor capable of producing the desired forces, a fairly strong magnetic field is required, and there is a need for a massive amount of back iron to carry the flux. Preliminary calculations indicated that a suitable linear motor would weigh at least 15 kg.

Several ideas involved using a rotary motor for mechanical power, and some sort of transmission to convert to linear motion. One concept involved a slider-crank four-bar linkage at the endpoint. Calculations quickly revealed that this idea was inferior to other transmissions, as the links would be fairly large to get the desired travel, and would require large torques and therefore a large motor.

A rack and pinion was considered for the transmission. The only real problem




Figure 2-1: Solid model of the new module for vertical motion.

was the bulkiness of the components needed to ensure structural rigidity with respect to the required forces.

A screw drive was considered and finally selected for the transmission. Power screws are typically fairly efficient, particularly if an effort is made to select components with minimum friction. They are also compact and available across a broad range of gear ratios. Calculations suggested that an acceptable tradeoff could be made between backdrive-ability and module mass by adjusting the travel per revolution (lead) of the screw (effectively changing the gear ratio).

2.3 Screw-Driven Design

The final design chosen for development consists of a screw-driven module mounted in place of the handle on a robot similar to MIT-MANUS. A solid model of the design is shown in Figure 2-1.

A servomotor is mounted at the top of the module and rotates the screw. The turning of the screw results in vertical motion of the nut, which is mounted to a linear bearing. The bearing system is designed so that the bearing takes the undesired side

forces and torques to which the nut is subject. The handle, the interface to the patient, is mounted to the nut. If the system is driven from the motor side, the turning of the motor produces vertical motion of the handle. If it is driven from the patient side, the movement of the handle by the patient turns the screw which in turn rotates the motor.

By far the most critical parameter of a design of this type is the screw lead. A smaller lead provides greater mechanical advantage in driving from the motor side, and therefore permits a smaller actuator and a smaller planar impedance. However, a shorter lead also increases the inertia and friction reflected to the endpoint, and increases impedance in the vertical direction. Based on a desire to minimize the reflected vertical inertia and friction as well as the planar inertia, a lead of 0.01905 m (0.75 in) per revolution was selected. Preliminary calculation suggested that the motor required for such a lead would result in a system with endpoint mass approximately equal in the horizontal plane and the vertical direction, and with reasonably low friction.

A more specific description of the key hardware components is found in the next chapter, along with a detailed description of how the system works in hardware and software.

Chapter 3

System Operation

This chapter provides a specific description of how the screw-driven therapy robot module works. The understanding that this chapter strives to impart to the reader is essential in interpreting the information provided in the remaining chapters, as all data presented was gathered on the system described here.

The first section of this chapter provides a discussion of the major hardware components, their essential features, and their interaction. This includes some discussion of modifications to the original design and suggestions for changes in future versions.

The second section discusses the key points relating to the structure of the software used to control the device. An understanding of the software proves critical in understanding the results in Chapter 4.

The final section of this chapter provides a simple description of how the system operates as a whole.

3.1 Hardware Overview

3.1.1 Rollnut

The most critical component of the module drivetrain is the screw and nut transmission. The lead was selected to be 0.01905 m (0.75 in) per revolution as described in the previous chapter. Given a specific lead, it became most essential to design

the screw package for minimum friction. The effect of greater friction on the forces required to drive the unit are more thoroughly discussed in Chapter 5, but it is fairly obvious from equation 2.3 that minimizing friction is always helpful in minimizing impedance.

Several designs for low-friction screw transmissions were considered. Primarily these included plastic lead screws, ball screws, and roller screws. Plastic lead screws rely on pure sliding contact and therefore offer fairly consistent performance. Special coatings are available to reduce friction, and efficiency numbers are often impressive. In fact, some manufacturers claim efficiencies as high as 50-80% for screw assemblies of approximately the correct dimensions [21].

Because they theoretically do not rely on sliding contact, ball and roller screws offer superior friction characteristics to even the best plastic lead screws. Ball screws are widely used in machinery but are subject to inconsistencies due to the rubbing of balls against one another. Many of the consistency problems can be removed by adding a preload, but this preload increases the friction force. Ball screw manufacturers routinely boast of efficiencies over 90% [13].

Roller screws are less common and generally more expensive than either plastic lead screws or ball screws. Several different designs are available; many of these still contain some sliding elements. Roller screws generally also offer extremely high efficiencies. Some roller screw manufacturers focus on their ability to create designs with extremely high load ratings and speed capabilities, especially as compared with ball screws [13]. In fact, much of the power screw industry is aimed toward these goals, which do not necessarily match the goals of a therapy robot.

Selection of the proper screw drive was aided through some realizations about the application. Unlike many machinery applications where power screws are typically used, this device does not require extremely tight position tolerances, nor does it require very large forces, or even particularly high speed. Instead the emphasis is on minimizing friction, even at the expense of position or force capabilities to some extent. Thus a screw that offers theoretically full rolling contact, with no sliding parts, is ideal.




Figure 3-1: Screw/nut interface of a Rollnut.

The prototype of the therapy robot uses a custom version of such a screw, called a Rollnut and manufactured by Norco, Inc. Figure 3-1 shows an illustration of the screw/nut interface. The nut is actually a carriage that holds two conical rollers. These rollers rotate freely in ball bearings about an axis perpendicular to that of the screw. The rollers ride in a precision track that is cut into the screw. The nut carriage has needle roller bearings at the top and bottom around the screw itself. The contact between these bearings and the screw is theoretically the only sliding contact as the nut moves up and down the screw, and this sliding is a minimal component of the motion. The Rollnut comes very close to offering pure rolling, and as such offers the lowest friction of any commercial screw product known to the author.

To have a package capable of supporting the desired load, a screw 15.875 mm (0.625 in) in diameter was selected. This system should be able to bear at least a 222 N (50 lb) load, clearly sufficient for this application. Figure 3-2 shows a photograph of the screw and nut package.

Though specific technical data about the roller package is unavailable, some measurements were made to determine the approximate dimensions. The outer shaft diameter is 15.875 mm (0.625 in). The inner shaft diameter, where the full thread




Figure 3-2: Rollnut assembly.

depth is removed, is about 12.7 mm (0.5 in). The width of the groove in which the roller rides is 3.0 mm (0.1185 in) at the thread ID and 4.0 mm (0.1575 in) at the thread OD. Thus if the roller is assumed to contact the screw primarily along the center of the thread surface, the effective screw diameter is 14.2875 mm (0.5625 in) and the effective roller diameter is 3.5 mm (0.1378 in).

3.1.2 Linear Bearing

For several reasons it is important to keep the screw and nut isolated from most side loads. The handle is connected to the nut, and it is desired that forces in the vertical direction be transmitted to the nut as completely as possible, while forces in other directions be transmitted to the structure of the SCARA linkage to effect the desired planar movements. It is important to recall that side forces are fully expected in this design, as they will be used to affect movement in the plane. Furthermore there are likely to be torques acting on the handle, as the environment is unpredictable. It is simple to see that a vertical force, applied at the handle 77.85 mm (3.065 in) away from the center of the nut, will produce both a force and a torque at the nut. This




Figure 3-3: A force at the handle induces a reaction force and a reaction torque at the nut and bearing.

simple 'beam' problem is illustrated in Figure 3-3. The distance from the center of the screw to the center of the handle acts as a moment arm and converts any forces that do not pass through the center of the nut into forces and torques at the nut.

If the nut is subjected to significant side loads and torques, this could be detrimental in terms of wear. Perhaps more significantly, substantial changes in loading could lead to changes in the friction forces in the nut through changes in the normal force. Thus the nut should be isolated from side loads and torques, while allowed to see the full vertical load.

The solution to this problem used in the design is a linear bearing. A BM15 rail and ball carriage system from Schneeberger was selected. The bearing is shown in Figure 3-4. It uses a recirculating ball design and includes wipers on the ends to keep the ball guidance system from becoming contaminated. The robot design is such that the bearing is placed as close as possible to the center of the nut. Thus any angular displacements at the bearing (any displacements will be extremely small) are not substantially amplified and do not allow any forces to reach the nut/screw interface. The center of the nut is only 14.351 mm (0.565 in) from the mounting surface of




Figure 3-4: Linear ball bearing, rail and carriage.

the bearing. The bearing has stated loading capacities of at least 9000 N in force and a minimum of 67 Nm in torsion. The loads to be used in this application are a miniscule fraction of these capacities, and should not affect the bearing at all.

3.1.3 Mounting Support

To transmit force from the bearing to the SCARA, a simple mounting support beam was designed. This beam was designed to be stiff enough to avoid significant bending even if 222 N (50 lb) of force was applied horizontally at the extremes of the screw. Part of the back of the support is cut away to reduce weight.

If the screw were to provide 0.46 m (18 in) of travel as originally planned, the support beam would have had to be extremely bulky and heavy. To keep weight down, it was decided to limit the travel of the screw to about 0.36 m (14 in). A photograph of the mounting support beam is shown in Figure 3-5.




Figure 3-5: Structural support/mount piece.

3.1.4 Flexible Coupling

The motor output shaft is connected to the screw shaft through a Helical Beam Coupling. This type of coupling is extremely stiff in torsion along the shaft axes, so as not to introduce resonance problems. The coupling is flexible with respect to torques about the other two axes, however, so as to alleviate the problems due to slight misalignment between the two shafts. A photograph of the coupling is shown in Figure 3-6. The coupling used can transmit up to 3.95 Nm (35 in×lb) of torque and is made of black anodized aluminum.

3.1.5 Alignment Difficulties

The screw system, as originally designed, is inherently overconstrained. The screw shaft is guided by a ball bearing at the top, mounted to a rigid bore, and by a needle roller bearing at the bottom, also mounted to a rigid bore. The ball bearing at the top locates the shaft in all three axes, while the needle roller bearing locates the shaft axis of rotation but allows play in the axial direction. The linear bearing attached to the nut results in a third point attempting to locate the rotational shaft axis. This




Figure 3-6: Helical beam coupling joining motor shaft to screw.

system is subject to classic problems of overconstraint.

Without extremely precise alignment, this overconstraint is likely to cause binding. During prototype assembly, repeated attempts were made to align the system carefully. Careful thought led to investigation of potential sources of binding. The potential problems, their sources, and efforts to evaluate the likelihood of each problem are shown in Figure 3-7. Measurements were made on a coordinate measuring machine, and the mountings were carefully shimmed to attempt to find the correct relative orientations of the parts. Figure 3.1.5 shows the top mounting bracket, that locates the top bearing bore, connected to the mounting support beam with shims in an attempt to achieve better alignment.

Despite repeated efforts to adequately align the system, binding was still an insurmountable problem. It also became obvious that if the system were ever disassembled, it would be extremely difficult to realign it to get it working again.

Because small errors in position are tolerable for this application, the possibility of adding some compliance to one or more of the bearings was appealing. It is disadvantageous to add compliance to the mounting of the nut to the linear bearing, as this forces the nut to bear side loads (see section on linear bearing). Instead it

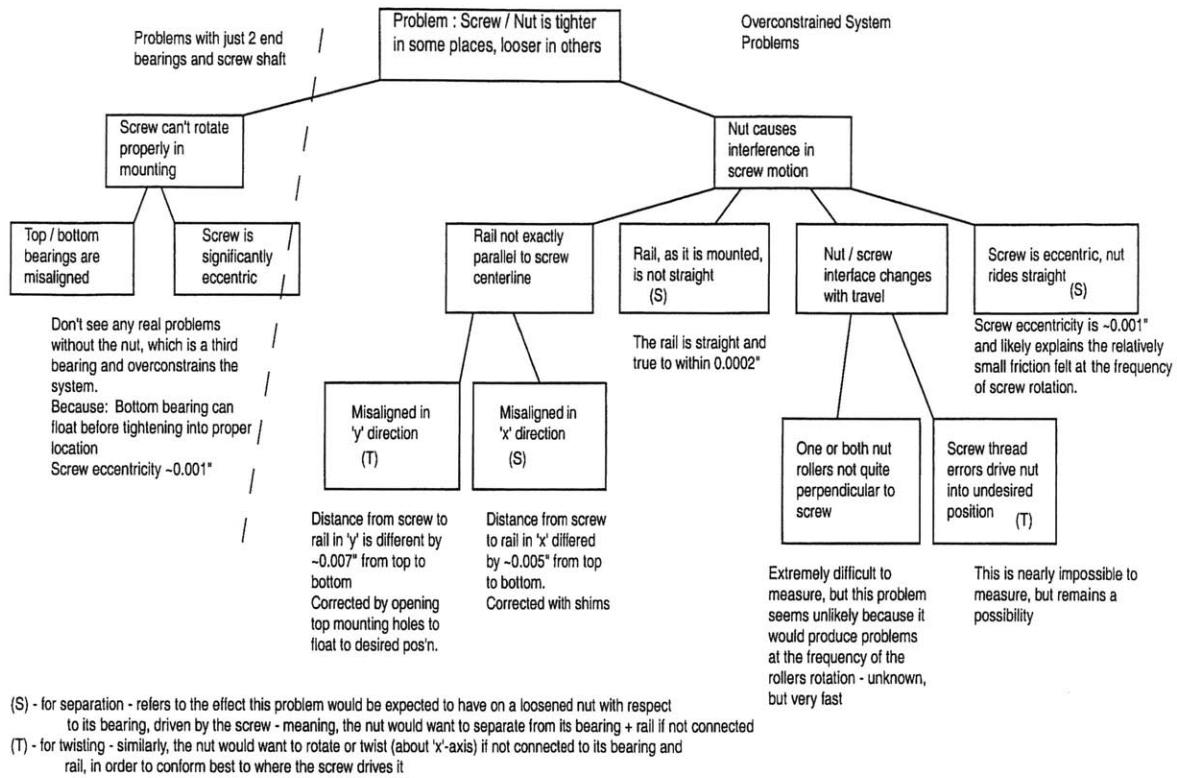



Figure 3-7: Summary of potential sources of binding in the screw assembly. Each potential problem area is investigated, and measurements are reported.






Figure 3-8: Grommet used to provide compliance in top bearing mounting.

was decided to introduce compliance at the top bearing mounting. Thus as the nut moves up and down, the screw essentially pivots about the bottom bearing into the position of least resistance, minimizing binding.

Introducing compliance to the assembly also introduces the risk of dynamic resonance problems. Appendix A contains an analysis of the risks of encountering such a resonance. The conclusion of this analysis is that it is important to maintain some minimum stiffness in the mounting to avoid the possibility of exciting a resonant mode with any of the driving frequencies of the system traveling at a reasonable speed.

The top bearing bore was enlarged to permit the addition of a compliant member. A rubber grommet like that shown in Figure 3-8 was added that fit snugly around the bearing and into the enlarged bore. This significantly alleviated the binding problems.

The idea of a compliant mounting seems like an excellent one for this design, and warrants further study for additional versions of the module. Perhaps compliance should be introduced at both the top and bottom bearings? More careful analysis should be done to select a proper compliance and to be sure of the actual compliance in use.

3.1.6 Servomotor

The system is powered by a brushless DC servomotor from Aerotech. The BMS-60 motor is a slotless three phase design for minimum cogging and breakaway torque. The maximum continuous torque T_{max} is 0.42 Nm (59 oz×in). The nominal torque constant is 0.16 Nm/A (22.7 oz×in/A) for peak current and 0.23 Nm/A (32.1 oz×in/A) for RMS current. The nominal rotor inertia is $1.96 \times 10^{-5} \text{ kg} \times \text{m}^2$ (0.0028 oz × in × sec²). With a screw lead of 19.05 mm (0.75 in) per revolution, the maximum endpoint force, assuming no losses, can be calculated with the expression:

$$F = \frac{2\pi}{\text{lead}} \times T_{max} \quad (3.1)$$

So the maximum continuous force is 138.5 N (31 lb). This is far enough above the desired value to allow substantial losses for friction.

Simple experiments were performed to test the breakaway torque and dynamic drag torque of the unpowered motor. The breakaway torque is defined as the torque required to start the rotor turning from a static position, in other words to overcome static friction and any cogging torque. The dynamic drag torque is defined as the minimum torque required to keep the rotor turning with the motor unpowered, in other words to overcome dynamic friction and cogging. Such constants can be useful in understanding the properties of a motor.

Weights were hung from a piece of fishing line wrapped around the output shaft of the motor. Applied torque was determined based on the weight and the radius of the shaft.

Breakaway torque can be dependent on position. Tests were repeated at several rotational locations. The average torque to break away was $5.44 \times 10^{-3} \text{ Nm}$ (0.77 oz×in). The dynamic drag torque was found to be approximately $4.98 \times 10^{-3} \text{ Nm}$ (0.705 oz×in).

3.1.7 Servoamplifier

The actuator receives power from a three phase Aerotech BAS-10 servoamplifier. The amplifier is capable of putting out from 0 to 10 A at a switching frequency of 20 kHz. The nominal bandwidth of the amplifier is 1 kHz.

For the experiments described here and the regular operation of the robot, the system is used in torque command mode. This means that a voltage commanded to the servoamplifier results directly in a corresponding current to the motor, and a corresponding torque output.

3.1.8 Position Sensor

The position sensor used for the system is an optical encoder that is part of the actuator package. The encoder, manufactured by Renco, has 5000 lines per revolution. There is an index bit that is triggered when the index location is passed to give an absolute position reference. The signal from the encoder is fed to the amplifier and is used for commutation. The signal is also sent to the computer, where a USDigital PC7266 encoder interface card receives it. This card interpolates the data and improves resolution by a factor of 4, giving a total of 20000 steps per revolution. Given the lead of the screw, this produces a position resolution of 9.5×10^{-7} m (3.75×10^{-5} in).

It should be noted that the sensor is collocated with the actuator, and therefore senses actuator position rather than true endpoint position. Thus any backlash or other play in the system will result in inaccuracies in data for actual position of the handle.

3.1.9 Force Sensor

Force of interaction with the patient is measured at the handle with a six-axis ATI force and torque transducer. The transducer has 16 bits of resolution and is calibrated to offer a resolution of 640 counts per lb (144 counts per N) for force and 640 counts per in×lb for torque (5665 counts per Nm).

Note that the force sensor is non-located with the actuator, and returns actual endpoint force and torque.

3.1.10 Adjustable Handle Module

The design for the module includes a mechanism to permit the adjustment of the angle of the handle for various moves. The idea is that the therapist can lock the handle into the position most suitable for the move before the move is made, so as to improve patient comfort and to ensure that the move is correctly trained.

3.1.11 Vertical Locking Device

The design includes a mechanism to physically lock the vertical direction in case only motion in the plane is desired. A pin is inserted into a hole at the base of the screw, preventing screw rotation.

3.2 Software Overview

The system is controlled by an extensive C++ library running in the QNX real time operating system. QNX is a UNIX- based operating system that is advantageous for real-time applications because it offers full access to priority scheduling. The software runs on a 550 MHz Pentium III PC.

The system is typically run at a sampling rate of 2 kHz. This is believed to be close to the maximum speed at which the code, as currently structured, can run. Table 3.1 gives an overview of the order of the major operations in the code in each sample period, along with some estimates of the approximate time after the start of the sample period that it occurs. As the clock signals the beginning of a new sample period, the software first reads analog input from the analog to digital (A/D) card, if there are any sensors registered that use those inputs. In the current configuration of this module, no analog inputs are used. Next the software reads position information from the encoder interface card. Then the force data is read. Necessary calculations

Event	Time (ms)	Approx. End Time
Clock Tick	0	0
Read position	0.050	0.050
Read forces	0.125	0.175
Make computations and call controller	0.100	0.275
Output command	0.050	0.325
Store Data/Graphics/etc.	0.150	0.475
Wait for Clock Tick	0.025	0.500

Table 3.1: Approximate timing of major software operations in each sample period.

are performed: the position is back-differenced to determine velocity (and this signal is filtered, if desired), and the control input is calculated. The control input is then sent to the digital to analog (D/A) card. Finally details such as data logging are taken care of, before the software finishes and waits for the next clock tick.

A more extensive timing chart, based on estimates from several experiments, is found in Appendix B.

3.3 Typical System Operation

A schematic of the system is shown in Figure 3-9. This figure is intended to show how each of the components described above is related so the system operation can be illustrated.

At each sample period, a desired position and velocity is specified in software. Measurements are made to determine the actual position and velocity. Since there is no velocity sensor in the system, this information is all gathered from the position sensor - the optical encoder - through the encoder interface card. The rotary position of the encoder is transformed into a vertical position for the handle. Velocity is computed by back-differencing the position, that is:

$$\dot{z} \cong \frac{(z - z_{prev})}{T_s} \quad (3.2)$$

where \dot{z} is an approximation of velocity, T_s is the sample period, z is the measured

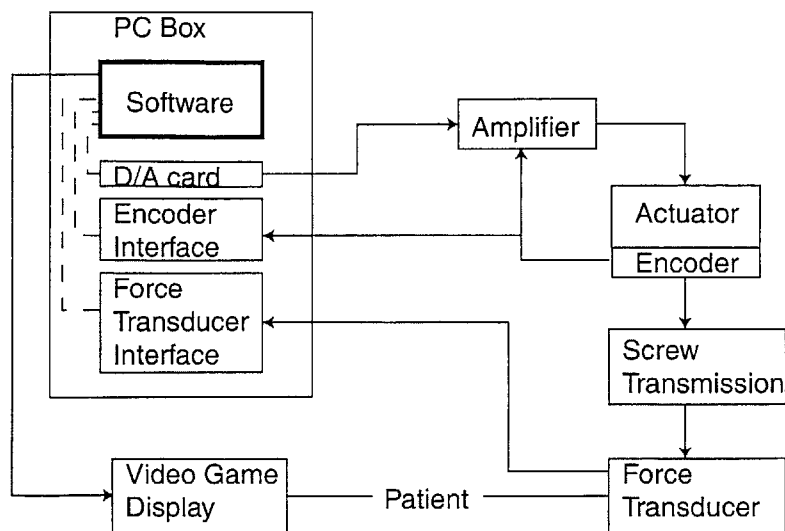


Figure 3-9: System schematic.

position, and z_{prev} is the measured position at the previous sample. When the quantity $(z - z_{prev})$ approaches the order of the position sensor resolution, this estimate becomes quite poor. The finite quantization of the position sensor and the small value of T_s result in a limited velocity resolution, and the signal picks up a substantial amount of noise at low speeds [12]. To smooth this signal, the velocity is then filtered with a first order Butterworth filter, with a cutoff frequency of 30 Hz [15].

Force data is also measured, recorded through the ATI transducer and transferred to software through the ISA card associated with the transducer.

In software the control law is applied; based on the actual position and velocity, and considering the desired position and velocity, a desired endpoint force is determined. For typical operation, the desired endpoint force will respond to some impedance law. Generally it is desired that the system feel as though there is a spring and damper between the desired location of the manipulator's endpoint and the actual location of the patient's hand [19]. This idea is illustrated in Figure 3-10. Following this idea leads to the basic structure of a proportional-derivative (PD) controller. Such a passive controller, when implemented in an ideal system, guarantees




Figure 3-10: A PD controller provides the feel of a spring and damper between the actual position and the desired position.

stability for interaction with any passive environment. The desired force is then:

$$F_d = -K(x - x_d) - B(\dot{x} - \dot{x}_d) \quad (3.3)$$

Equation 3.3 could be used as the control law and, for an idealized system with zero damping and stiffness, the actual endpoint force would equal F_d . MANUS, for example, is very close to such an idealized system within the workspace. Equation 3.3 is used as a control law for MANUS, and performance approaches desired levels [15]. For the new screw module, friction is a substantial limitation, and a more advanced controller may be needed to produce the desired endpoint force. The actual control law used may include some effort to compensate for the friction and inertia in the drivetrain, and may include terms other than those shown in equation 3.3. For example, the measured force may be incorporated into the control law. This idea will be discussed in greater detail in subsequent chapters.

Once the control force input is determined, it is converted to a torque and then to a current. The D/A card is then commanded to output this value to the servoamplifier. The amplifier produces the commanded current and sends it to the motor. The motor converts the current to torque, and the screw transmission converts the torque to

force. This force is then transmitted to the environment (the patient).

The above process is repeated 2000 times per second, once for every sample.

Integration with MANUS

Recall that the module, largely discussed here as an independent machine, will be integrated with a planar SCARA robot to form a complete three-degree of freedom therapy robot. The process described in the preceding paragraphs is very similar to that used in MANUS. The two systems, in fact, already run very similar software libraries.

When the modules are integrated, they will run under one piece of software on a single PC. The control algorithm for the vertical motion will remain decoupled from that for planar motion. Each of the cards associated with the sensors is expandable and can accommodate all three degrees of freedom. A second D/A card will be needed to provide three channels of output.

The integration of the two systems is a straightforward task that is well understood. It is important to keep in mind that the system will at one time be integrated, to make sure that there will be no interference between the processes the two systems need to execute.

3.4 Studying the system

For purposes of the study of the system detailed in the rest of this thesis, the system operations can be broken up into several major processes. These processes are illustrated in block diagram form in Figure 3-11.

The first block is the software. This block has been described above, and more information is provided in Appendix B. Note that the software accepts sensor information and program parameters as inputs, and produces a voltage representing the desired torque as an output. All these inputs and outputs are signals, and are therefore shown with dashed lines.

The next block is the subsystem converting the signal and electrical power from the




Figure 3-11: Block diagram of major subsystems for the therapy robot module. Dotted lines denote signals. Sets of solid lines mark power exchange.

wall to mechanical power. This includes the cards in the PC, the servoamplifier, and the servomotor. The input is the desired torque, and the power output is an actual torque transmitted through a rotation. This block is analyzed and characterized in Chapter 4.

The final block is the mechanical transmission, that converts power in rotation to the endpoint power interaction with the patient. This includes the screw transmission, the flexible coupling, and the linear bearing. This block demands substantial study because it contains the friction and inertia parameters that are so important, as well as the influence of gravitational forces on the device. Note that the hardware cannot be divided so as to fit in these blocks in the strictest sense; the servomotor plays an important role in each of the power transmission blocks, as drawn here. The mechanical transmission block is studied in Chapter 5.

Fully understanding the blocks in Figure 3-11 and the interactions between the blocks provides the information needed to design a controller for stable interaction and desired performance. It is nearly impossible to gain a complete understanding because of the system's complexity and its susceptibility to slight changes over its life. It is, however, possible and desirable to obtain a level of understanding that is commensurate with the desired level of control, and to bound uncertainty such that stability and performance, within some bounds, can be guaranteed. The effort to reach this understanding comprises the bulk of the work described here.

Chapter 4

Signal to Mechanical Power Subsystem

4.1 Amplifier and Actuator analysis

4.1.1 Low Frequency Static Testing

As described in Chapter 3, the brushless DC servomotor is operated with its servoamplifier in torque command mode. A voltage between -10 V and +10 V is commanded to the servoamplifier, which theoretically delivers three phase current with the commanded magnitude in amperes, between -10 A and +10 A. This current is converted to torque by the motor. The relationship between current and torque is theoretically linear and is determined by the torque constant (K_t) of the motor. The BMS-60 used in the robot has a published torque constant of 0.16 Nm/A (22.7 oz×in/A) for peak current. Thus a voltage command of +1 V should produce +1 A from the servoamplifier, and +0.16 Nm of torque from the motor.

An experiment was designed to measure the actual torque constant of the motor and to check for linearity of torque with current. The motor was mounted to ground with its rotor locked to ground through the ATI force transducer, as shown in Figure 4-1. Thus the rotor was prevented from rotating significantly. An input sinusoid was commanded from software at a frequency of 1 rad/sec (0.16 Hz) for approximately




Figure 4-1: Locked rotor testing. The servomotor is on the left, and the force transducer on the right.

7 seconds, so that more than one period of oscillation was completed. Torque was measured with the force transducer and recorded in software. Figure 4-2 is a plot of the output torque versus the commanded current. Clearly the torque is nonlinear close to the origin; this nonlinearity impacts torques between roughly -0.025 Nm and $+0.025$ Nm.

Further investigation with apparatus to be described in the following section showed that the nonlinearity was present in the actual current as well as the torque. The servoamplifier was identified as the source of this nonlinear behavior. This problem is likely a result of the switching bridge circuits in the servoamplifier output stage. Apparently the output of these bridge circuits has not been linearized, and it does not follow the desired pattern close to zero command.

Since the system is to be used with a linear impedance controller, it is desirable to linearize the torque/current behavior. This was done in software by mapping the desired torque to a current that will produce this torque. To design this part of the software, the (command signal)/(actual torque) curve was fitted. This was best accomplished with a piecewise fit; the curve was split into three sections, and the top

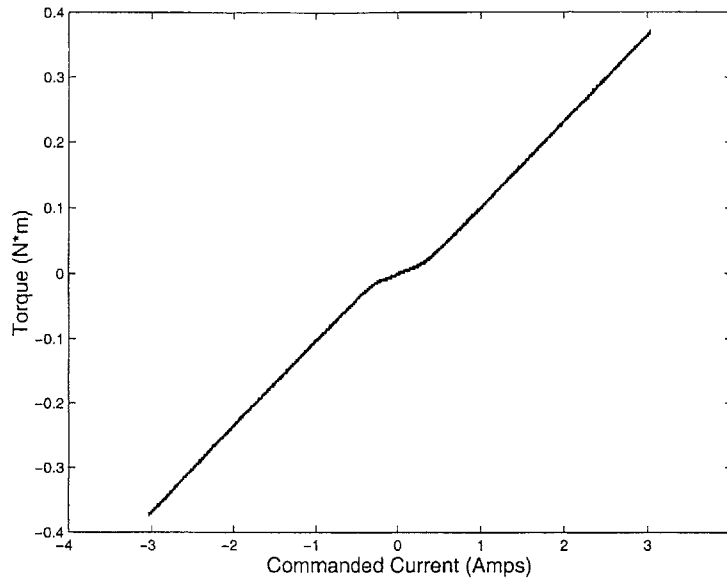


Figure 4-2: Torque versus commanded current, BAS-10 servoamplifier and BMS-60 servomotor.

and bottom segments each were fit with a line, using the least squares criterion. The middle segment, from desired torques of approximately -0.025 Nm to $+0.025$ Nm, was fit with a third order polynomial. This curve is shown in Figure 4-3. Note that the function is forced to be continuous. The first derivative is discontinuous at each junction of the curve segments, but the jump is on the order of 10 percent of the derivative at that value. This curve is sufficiently smooth for the application.

The new mapping for desired torque was implemented in the software, and the experiment was repeated. The result is shown in Figure 4-4. This curve is close to linear across the entire torque range, with a correlation coefficient of 0.99996. This adjustment satisfactorily linearizes the amplifier performance. Further testing with the robot library software used the linearized version unless otherwise noted.

4.1.2 Amplifier Frequency Response

To fully understand the behavior of the amplifier block in Figure 3-11 it was necessary to observe any significant dynamics of the amplifier with frequency response testing. It is expected that amplifier electrical dynamics should only be prominent at such high




Figure 4-3: Commanded current versus desired torque. This piecewise function compensates for the nonlinearity shown in Figure 4-2.




Figure 4-4: Torque versus current, using linearization.

frequencies that they have little effect on the mechanical domain of the interactive robot. Prior work has shown, however, that three-phase switching amplifiers can show evidence of pure delay [15]. Pure delay has proven to be a serious limiting factor for stability, and should therefore be modeled and understood.

Amplifier frequency response testing was done in two phases: First the amplifier was tested with a Dynamic Signal Analyzer over a broad range of frequencies to identify any delays and dynamics. Then the system was tested with the discrete time sampling of the robot libraries in an attempt to identify software delays and to clearly model the z-domain behavior expected when actually running the system in QNX.

Testing with Dynamic Signal Analyzer

The frequency response of the servoamplifier was tested with a Hewlett Packard 3562A Dynamic Signal Analyzer (DSA). The DSA has the capability of inputting to a system sinewaves across a broad range of frequencies, recording the system input and output, and generating a frequency response plot of magnitude and phase in Bode form. The DSA samples at 256 kHz and then filters data based on the frequency at which the system is driven. Note that a sampling frequency on the order of 256 kHz results in an apparent delay due to sampling of less than 2 microseconds. Given this fact, the data from the DSA is treated here as continuous time data.

The DSA was used to perform a swept sine analysis on the amplifier. The frequency was swept from 1 Hz to 10 kHz, with 40 frequency points per decade. For the system input signal, the DSA output was connected to the command input on the servoamplifier. The current from one of the phases (phase C) of the amplifier/motor was passed through a 0.1Ω resistor. Voltage was read across this resistor and passed through a differential amplifier (Burr-Brown INA117). The output of this differential amplifier was connected to the DSA input, completing the circuit.

The results of testing with the DSA are presented in Figure 4-5. Clearly the circuit shows a rolloff at around 1 kHz. This data nicely fits a first order Butterworth or Bessel filter. The same data fitted to a model is shown in Figure 4-6. The model used is a first order Butterworth filter with a cutoff frequency of 1 kHz, plus a 60




Figure 4-5: Frequency response data taken with Dynamic Signal Analyzer: Magnitude and Phase versus frequency, Current/Commanded Current.

μ sec pure delay. This appears to be an accurate characterization of the amplifier.

Testing in the Libraries

In addition to understanding the dynamics of the amplifier, more can be learned by observing its frequency response when run in the software libraries. This was accomplished by using the robot library software to command an input sinusoid to the servoamplifier, and to record the current. In order to read the current with the A/D card, it was necessary to use an existing circuit to low-pass filter the signal and then to amplify it to cover more of the available range of the A/D board. This current sensor circuit reads voltage across a 0.1Ω resistor through which the current flows. This voltage is input to a Burr-Brown INA117 Differential Amplifier. The output is passed through a low pass filter using Analog Devices OP177 operational amplifier. Circuit parameters were selected to provide a cutoff frequency of approximately 1 kHz (928.8 Hz) for this filter. The output is then input to a gain amplifier created with another Analog Devices OP177. This gain makes 1 A of current through the resistor theoretically appear as 3.42 V at DC. All three channels of the current sensor




Figure 4-6: Frequency Response data as in Figure 4-5 with fitted model.

were calibrated. The sensor dynamics must also be understood in this analysis, as some phase lag may be contributed from the filter portion of this circuit.

The frequency was input in 40 logarithmically spaced steps ranging from 1 Hz to 200 Hz. The output sinusoids were recorded, then fitted with a Matlab program to determine magnitude and phase of the three phase signal. Unfortunately there was a 60 Hz noise signal present on top of the data. Because the entire circuit was exposed to the nearby power cables of the system power electronics, it proved difficult to eliminate this noise. Since the data was only to be examined offline and not in real time, it was decided to use a discrete time band-stop filter to remove frequency components around 60 Hz. An optimal filter was implemented using `remez` in Matlab. The data around 60 Hz is somewhat obscured by this processing, but the rest of the data remains unaffected.

The results of the testing and data processing are presented in Bode form in Figure 4-7. A model including the known amplifier dynamics, the known circuit dynamics, and the known amplifier delay, was compared to the data. The data showed additional phase lag. This is not surprising, since some software delay between 100 and 400 μsec is expected, as described in Chapter 3. This delay represents the computation time




Figure 4-7: Frequency response data sampled at 2 kHz: Magnitude and Phase versus frequency, Current/Commanded Current.

between reading data from the A/D and writing on the D/A for the same sample.

The modeled software delay was adjusted by hand until the model fit both the magnitude and phase frequency response. The data and fitted model are shown in Figure 4-8. The modeled software delay was 150 μ sec. This is slightly lower than expected. This will be discussed in more detail in the following sections. The model used in the simulations that produced this fit is discussed in Appendix C.

4.1.3 Torque frequency response

An understanding of the behavior of the actuator in response to an input current is critical for making the robot behave as desired, producing the desired endpoint force and maintaining stable performance. To understand the actuator dynamics a frequency response test was performed. The experiment was designed to measure the output torque resulting from a commanded input current.

The test was performed with the motor rotor locked with respect to the frame, so no substantial rotation was permitted. The rotor was connected to ground through the ATI six-axis force transducer, so that the torque generated by the motor was




Figure 4-8: Frequency response data as in Figure 4-7 with fitted model.

detected as torque in the z-axis of the transducer. The input command was delivered through the servoamplifier; a sinusoidally varying voltage was generated by software and the D/A card, and input to the servoamplifier, which in turn produced a corresponding current. Note that the dynamics and delays due to the amplifier, the sampling, and the software are also expected to be present in the data from this experiment. The test setup is shown in Figure 4-1, the same as for the prior experiment.

Current was input to the motor at 40 different frequencies from 1 Hz to 200 Hz, spaced equally on a logscale. The input amplitude was approximately 2.2 A. A sinewave was fitted to the measured torque data at each frequency. With no motion of the rotor, the effects of the motor damping and inertia are irrelevant, and the motor can be modeled as a constant multiplier as follows:

$$T_m = K_t \times I \quad (4.1)$$

In the above equation K_t is the motor torque constant, in Nm/A, and T_m is the motor torque.

The data from the torque frequency response is presented in Bode form in Figure 4-

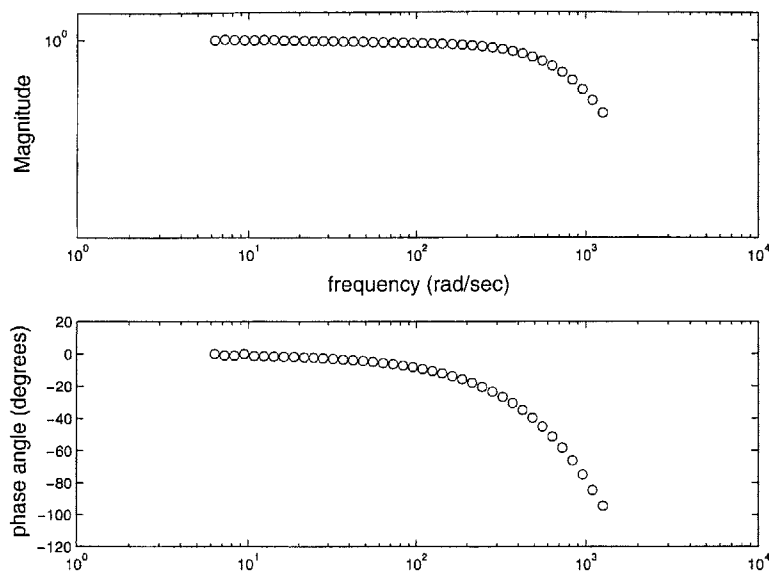


Figure 4-9: Frequency response data sampled at 2 kHz: Magnitude and Phase versus frequency, Torque/Commanded Current.

9. Note that both the magnitude and phase show substantial rolloff within the tested frequency range. This was not expected, because the mechanical actuator dynamics have been removed (by locking the rotor) and the electrical dynamics should only appear at much higher frequencies.

An aliased resonance

The test setup was examined to determine the source of the magnitude and phase drop at low frequencies. As noted previously, the Aerotech BMS-60 motor has a published inertia of $1.96 \times 10^{-5} \text{ kg} \times \text{m}^2$ ($0.0028 \text{ oz} \times \text{in} \times \text{sec}^2$). The ATI force transducer has a published stiffness with respect to torque in the z-axis of 10,000 Nm/rad. This leads to an undamped resonant frequency of 3600 Hz.

Recall that in this experiment the system is sampled at 2 kHz. Thus the Nyquist frequency is 1 kHz. Due to frequency folding, high frequency effects are likely to appear at frequencies below the Nyquist frequency. In this case, the undamped resonance described in the previous paragraph would then appear at 400 Hz, 1600 Hz,




Figure 4-10: Frequency response data as in Figure 4-9 with fitted model.

3600 Hz, and so on, according to the formula

$$w_a = n \times w_s \pm w_e \quad (4.2)$$

where w_a is the set of all apparent frequencies at which an event that is actually at frequency w_e appears. n is any integer. The predicted frequency of 400 Hz does not quite match the frequency observed in the data, but this can be explained as well.

Since the rotor is moving rather than fixed, the testing is no longer truly locked rotor and the damping cannot be neglected. The actual damping on the mode is a combination of friction in the motor bearings, structural damping in the force transducer, and damping due to the motor resistance. The model applied to the data was modified to include this mode. This model included the amplifier dynamics and sampling effects from the prior testing. An additional pure delay term was also included in the model to represent the software delay, and the damping and delay were adjusted to fit the data. The data and model are shown in Figure 4-10. The model used to generate this fit is provided in Appendix C.

A linear plot of phase versus frequency is shown in Figure 4-11 and emphasizes




Figure 4-11: Torque frequency response phase data with model, linear scale.

the closeness of the model to the data. Note that the model fits quite well with a damping value of 7.2 Nm/rad/sec and a software pure delay of 400 μ sec. This is greater than the delay detected in the analysis of the current data. This discrepancy will be discussed in the following section.

To increase confidence in the several assumptions made in analyzing the data, the torque frequency response test was repeated with a sampling rate of 1 kHz. The model fit the data to a satisfactory degree at frequencies below about 100 Hz. At higher frequencies the fit was not as good. Note, however, that in these frequency ranges a 1 kHz sampling rate provides less than 10 samples per period, and it is difficult to fit the data well. Simple examination of the data reveals that it is not, in fact, accurately fit at high frequencies. Data sampled at 1 kHz is provided in Figure 4-12. The model was deemed acceptable, except for the discrepancy in the software delay from the two sets of tests.




Figure 4-12: Frequency response data sampled at 1 kHz: Magnitude and Phase versus frequency, Torque/Commanded Current, with fitted model.

4.1.4 Limitations of delay information

Recall that the model fit to the current data suggests a software delay of 150 μsec , while the model fit to the torque data suggests a delay on the order of 400 μsec . Direct study of the software, though not completely rigorous, suggests a delay somewhere between 200 and 450 μsec .

Much as Nyquist's theorem says that information beyond a certain frequency (half the sampling frequency) is unavailable in sampled data, there is a similar limit on the observable effects of pure delay in sampled data systems. Pure delay adds a linear increase in phase lag with frequency, and thus can be modeled in a dynamic system sampled discretely by subtracting the appropriate angle from the phase angle suggested by the dynamics. In the experiments described above, this modeling technique was used and the delay was adjusted until the model aligned well with the data.

Consider an ideally sampled system where the output/input relationship is simply a unity gain. One can imagine simply passing current through a wire with zero resistance and measuring the input and output. The fact that the system is ideally

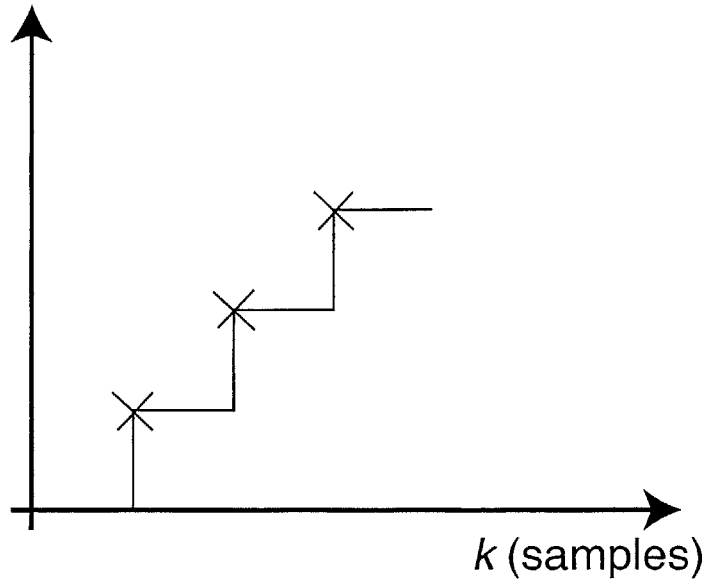


Figure 4-13: Ideal sampling of a unity gain plant.

sampled means that the read and write theoretically occur at the same instant [16]. Since the system is discretely sampled, the input is a series of small steps. This situation is illustrated in Figure 4-13. The solid line is the input and the output, since the two signals are identical. The x's indicate the read output data. When these x's are connected with a zero order hold, the output is reconstructed to exactly match the input.

Consider now the case shown in Figure 4-14, where the system has a delay T_d such that

$$0 < T_d < T_s \quad (4.3)$$

where T_s is the sample period. The solid line is the actual input, discretely input using a zero order hold. The dashed line is the actual output, a delayed version of the input. The x's again represent the sampled data points. The dotted line represents the reconstructed output data from connecting the x's with a zero-order hold. Note that regardless of the value of T_d , the reconstructed signal will appear exactly one




Figure 4-14: Sampling a system with a partial sample delay.

sample late. It is easy to extend this idea to see that if

$$T_s < T_d < 2T_s \quad (4.4)$$

the reconstructed signal will appear 2 samples late. Thus it is impossible to resolve delays with any precision finer than one sample period in looking at sampled data.

For the testing done here, it is best to assume a 500 μsec delay from software. This is a sufficiently conservative assumption because it is not possible for the delay to be any larger and have the software run at 2 kHz.

4.2 Subsystem conclusions

The amplifier and actuator characterization provides insight into the workings of the sampled system as a whole and provides validation for a number of modeling assumptions. The pure delays in both the amplifier and the software are significant and are modeled in future system analysis. The amplifier nonlinearity was effectively removed in software. The first order filter in the amplifier could be significant and

will also be modeled.

The actuator testing makes it clear that the actuator has no significant dynamics in the expected system operating range. The electrical dynamics can be neglected and the motor can be modeled as a damped mass. The published rotor inertia value appears to be correct. The damping was not directly determined, as the motor damping could not be decoupled from possible transducer damping in this experiment. This is an insignificant weakness however, as the system friction will be examined extensively in the next chapter, and the motor friction will be included in this analysis.

The subsystem providing conversion from a software command to mechanical power is sufficiently understood from a system modeling point of view.

Chapter 5

Transmission Subsystem

This chapter discusses the characterization of the transmission of the therapy robot module, and the identification of relevant parameters to describe this transmission. The components considered can loosely be grouped into the box labeled “mechanical transmission” in Figure 3-11.

It should be noted that although the actuator was characterized as part of the previous chapter (and as part of the previous block in Figure 3-11), it is generally included here as well. In the prior chapter, the actuator was considered as a non-moving torque source. Mechanical dynamics of the motor were neglected except when needed to explain unexpected results.

In this chapter, parameters such as rotor inertia and motor friction are considered part of the system being tested, except where otherwise noted. Generally the rotor inertia is combined with the screw and coupling inertias to form one lumped inertia value, and the motor friction is considered as a contributor, along with several other components, to the total system friction.

5.1 Parameters to be identified

As has been mentioned in prior chapters, the system needs to be driven by both sides; the motor drives the screw to pull the patient up and down, and the patient pushes the handle and turns the rotor. These two functions must be performed, at

different times, either independently or together. For example, with a totally flaccid patient in early therapy, the robot must be able to move the load of the patient's arm (which may be relatively constant or time varying) from point to point. In evaluation sessions or later therapy, the robot may be effectively turned off, and the patient must be able to push the handle around. In intermediate stages of recovery, the robot could mainly provide corrective action to the patient's movements. Clearly, the module must accommodate power input from both sides.

As a transmission from rotation to translation, the subsystem discussed here has both parts that translate, such as the nut, handle, and linear bearing carriage, and parts that rotate, such as the screw, the motor rotor, and the rollers in the nut. Parameters such as friction and mass or inertia originate in the translating parts as well as the rotating parts. Despite this duality, the system still only has one degree of freedom. For simplicity, the system can be viewed from either port as a simple moving mass; from the rotational frame, it is simply a rotating inertia, and from the translational frame, it is simply a sliding mass. Parameters can be lumped together and considered in either frame. Transformations can easily be made between the two frames, by simply considering the system geometry. Translation z can be related to screw rotation θ by the lead of the screw L :

$$z = L \times \frac{\theta}{2\pi} \quad (5.1)$$

Similarly, torques and forces are related by:

$$F = \frac{2\pi}{L} \times T \quad (5.2)$$

Because of the special interest in driving the system from the patient's side, the system is considered in the translational frame. The equation of motion for the system can be given as:

$$M\ddot{z} + F_f(z, \dot{z}) + mg = F_c + F_p \quad (5.3)$$

where z is vertical position, M is the effective endpoint mass, F_f is the friction force




Figure 5-1: Free body diagram of the equivalent system in translational space.

opposing motion as a function of position and velocity, mg is the force due to gravity, F_c is the control input force and F_p is the environmental force, from the patient. This equation is trivially derived by applying Newton's second law of motion to the free body diagram shown in Figure 5-1.

Note that equation 5.3 can be easily converted into the rotational equation by substituting equations 5.1 and 5.2 as well as the following expression for converting mass to inertia:

$$M = I_{eq} \times \left(\frac{2\pi}{L}\right)^2 \quad (5.4)$$

The only difficulty in this transformation comes in dealing with the friction force because it is a nonlinear function of z as well as \dot{z} . When the friction force is known for a given state, however, it can be easily translated with 5.2. The resulting rotational equation of motion has the same form:

$$I_{eq}\ddot{\theta} + T_f(\theta, \dot{\theta}) + T_{mg} = T_c + T_p \quad (5.5)$$

where each parameter is a rotational equivalent of the translational terms in equa-

tion 5.3. This conversion is provided only to justify that working in translational coordinates does not neglect anything.

Examination of equation 5.3 reveals several components that can be easily determined, and several that cannot. F_c , the control force, is presumably always known to a fairly close precision, as it is a commanded value. F_p can be measured, as can z . At least in theory, \dot{z} and \ddot{z} can be determined numerically with a time history of z . Three major elements of equation 5.3 are left unknown: M , F_f , and mg . Note that mg is not simply the equivalent endpoint mass M multiplied by the gravitational constant; because part of M is actually a reflected inertia that does not gain potential energy as z increases, the mass that leads to mg is actually considerably smaller than M and represents only parts that actually move up and down.

5.1.1 The importance of friction

Perhaps the most daunting and important unknown quantity is the friction force. F_f is assumed to be a function, potentially a nonlinear function, of position and velocity. Were it not for the potential nonlinearity of F_f , the identification of the other parameters would be considerably easier, as equation 5.3 would reduce to a linear equation.

Furthermore it becomes clear from just a few minutes of handling the screw module that friction is a dominant factor in the performance of the design. From a very qualitative perspective, it is obvious that friction varies significantly with position. In pushing the handle from top to bottom, one can feel regions of substantial binding, and variations periodic with the rotation of the screw are apparent as well.

Recall that the system was carefully aligned and compliance was added at the top bearing to reduce binding. The binding was minimized to a reasonable degree, but it is still present. So it is important to note that any identification and characterization of the friction in the device is only valid for this specific prototype, and no longer applies if the device is disassembled.

Friction also presents a tantalizing problem to be addressed. Recall from Chapter 2 that minimizing friction was of utmost importance in the design process, and that

sacrifices were made in inertia in order to minimize friction. In fact, many common machine elements have a substantial amount of friction. If the friction can be well understood and its effects can be minimized, the door could be open to future designs with more friction, leading to gains in other areas. This particular interest in friction led to the dedication of a significant portion of this work to its understanding.

5.1.2 Why do we need to know?

Before detailing the process by which the system parameters were identified, it is important to clarify the potential rewards for such a significant undertaking. As noted in Chapters 2 and 3 the system is intended to be used with an impedance controller, to present the feeling of a virtual spring and damper to the user. Such a controller will permit stable interaction with a large variation in environmental conditions. While a passive control algorithm such as a PD controller, with any nonnegative gains, interacts stably with any passive environment, the control law is actually implemented in a nonideal system. This system includes sensors with finite resolution and potential for measurement error, and discrete time sampling with a finite sampling rate. Thus to ensure stability (an absolutely essential requirement for this application), it is necessary to model the system. Without accurate knowledge of the parameters, it is impossible to have much confidence in system models.

For this application characterization is of special interest. The system performance at low impedance is likely to be unsatisfactory with a simple PD controller. This is largely due to the size of M and the significance of F_f in equation 5.3. Alternative control algorithms may be explored to improve performance, particularly at the low end of the impedance spectrum. Having an accurate model of the system permits useful simulation in the controller design process, and opens the door to model-based control methods. Control will be further discussed in the next chapter.

System characterization and modeling are valuable steps in the design process. As this is a prototype, specific data gathered about this device can be used to improve the design for later versions.

5.2 Friction Identification

The process of identifying F_f in equation 5.3 involved first studying friction in machine elements in the literature to understand the phenomena to expect and the physics behind these phenomena. Then a series of experiments was carried out to gather data about the system at hand. The results of these experiments were correlated to the behavior predicted by theory.

5.2.1 Understanding Machine Friction

Friction is one of the greatest limitations to performance in precision machinery. Friction affects everything from power consumption and machine life to positioning accuracy and control stability, and thus has a huge impact on design.

Despite its importance, friction remains, in many ways, a bit of a mystery. This lingering misunderstanding certainly does not come, however, from a lack of effort. A voluminous body of literature has been devoted to tribology, the study of friction. The issue is approached experimentally, through modeling and simulation, and in theory at the molecular level. Much of the tribology literature draws on material science and fluid mechanics. Friction is such a complex phenomenon that it generally is studied in extremely focused terms, with simple and highly specialized experiments and simulations.

The daunting problem of friction must be simplified to some degree in order to study it in the context of machinery. Understanding all the friction in a complex machine to a precise, molecular level would require an unrealistic amount of time and effort. Several common phenomena tend to dominate machine friction. Generally machines involve some combination of parts that roll and parts that slide on each other, almost always with at least some form of fluid lubrication in one or more interfaces. These types of contacts thus earn primary focus here.

Sliding and Rolling

Both sliding and rolling friction exhibit many of the same fundamental characteristics. While friction force between bodies in rolling contact is generally far less than that between bodies in sliding contact, Armstrong et al. point out that the primary mechanism producing friction in members thought to be in rolling contact is, in fact, sliding [4]. Consider a ball rolling in a raceway, as often found in a ball bearing. Since both the ball and the raceway deform elastically, the contact between the two bodies is not a simple line, but is actually an ellipse. Within this ellipse, two curvilinear segments are at any given time in rolling contact, but the rest of the contact area involves sliding, either in the direction of rolling or in the opposite direction. Since the fundamental mechanism is the same, it is acceptable to consider rolling contacts as if they were sliding contacts. This assumption is supported by the work of Dahl and its use in the tribology community [9, 10]; Dahl's original work was done with ball bearings, involving theoretically rolling parts, but has since often been extended to sliding parts.

The Physics of Friction

The field of tribology has achieved a strong qualitative understanding of the mechanisms that produce the range of behaviors seen in the interaction of contacting solids. The actual physical contact of two bodies covers a much smaller surface area than the perceived contact area. This is because the two bodies actually only touch where the peaks of a number of small protuberances on each body, known as asperities, come into contact (e.g. [3, 4, 6]). This situation is depicted in Figure 5-2. For the bodies to move relative to each other, it is necessary either for the asperities to deflect or for the contact between them to be broken. Each of the types of motion to be discussed in subsequent paragraphs fall into one of these two categories.

Basic physics and engineering texts often teach of several simple types of friction. These ideas form a solid foundation for understanding more complex concepts about friction. Viscous damping, or hydrodynamic drag is common for fluid interfaces to



Figure 5-2: Contact between two solid bodies is limited to the contact of small protuberances called asperities.

solid objects. Viscous damping results in a friction force that is proportional to the relative velocity between the solid and the fluid, following the relationship

$$F_v = -b \times v \quad (5.6)$$

where F_v is the friction force due to viscous damping, v is the relative velocity, and b is the damping coefficient. Note that F_v acts in the direction opposite to v . Viscous damping is a good approximation for many friction systems involving fluid contact, and is especially convenient for dynamic analysis because it is linear.

Sliding contact between solid bodies is often approximated with Coulomb friction and stiction. Coulomb friction results in a friction force F_k proportional to the force F_n normal to the direction of motion:

$$F_k = \mu_k \times F_n \quad (5.7)$$

μ_k is the coefficient of kinetic friction. Note that equation 5.7 is independent of the relative velocity. If the relative velocity between the two bodies is zero, however, the force to produce initial relative motion is often a bit higher than that predicted by equation 5.7. The static force F_s is defined by:

$$F_s = \mu_s \times F_n \quad (5.8)$$

Note that typically, $\mu_s > \mu_k$. Thus a force greater than F_s , known as the breakaway




Figure 5-3: Friction force versus relative velocity for two solids in lubricated contact.

force or the force to break stiction, is required to produce relative motion, and a force equal to F_k is needed to maintain a constant relative velocity.

The dependence of the friction force on the relative velocity of the two bodies is well recognized and often studied, most trivially in the context of viscous damping. A typical relationship between the relative velocity of two bodies and the force in the direction of motion needed to maintain this velocity at steady state is presented in Figure 5-3. This figure is divided into several regimes indicating different characteristic behavior [4], and includes several characteristics more advanced than a simple study of Coulomb and viscous friction. At high velocities, the curve represents a combination of these simple approximations; it is clearly the sum of a nonzero force due to Coulomb friction, constant with velocity, and a component linear with respect to velocity. At zero velocity, the force is at a local maximum, exhibiting stiction behavior. The rest of this curve is addressed in the coming paragraphs.

Before the asperities have fully come out of contact with one another, in other words before the slip region is entered, the friction force is dependent on position rather than velocity. The understanding of this region is attributed largely to Dahl, who found that in this region the contacting asperities deflect and act like a spring.

Dahl found the relationship between position and force to be close to that of a linear spring, with some small amount of structural damping [9, 10]. The Dahl effect, as it is known, is an important point of understanding, because if only velocity dependence is considered, there is a large range of possible forces at zero velocity ranging from the breakaway force in the positive direction to the breakaway force in the negative direction.

When the slip region is reached, the “springs” of the Dahl effect have essentially broken, and friction force becomes dependent on velocity. The first region of this type involves slow motion against a solid boundary layer [4]. Velocity is not sufficient to draw fluid lubricant into the contact region, so two solids are sliding against one another. Friction force is lower than the breakaway force, but still fairly high. Depending on the system, this regime may only cover a tiny range of velocities, and therefore may not be very obvious.

The behavior found at the next range of velocities is extremely important for understanding machine friction and its associated problems. This is the region of partial fluid lubrication. Note the downward trend in force with increasing velocity. This behavior is known as the Stribeck effect and is known to wreak havoc on efforts to stably control a system with significant friction. In this regime, the velocity becomes great enough to begin drawing fluid lubricant into the contact area. As more and more fluid is drawn in, the friction is reduced as the fluid interface reduces the area of contact between actual sliding solid bodies. The velocity at the point where the slope of the force curve becomes zero is known as the Stribeck velocity [4].

The Stribeck effect amounts to a region of negative viscous damping - close to a linear relationship between velocity and force, but with a negative damping coefficient. This results in substantial stability problems for controllers in systems with friction. An interpretation like the Stribeck effect is needed to allow a stiction force that is higher than a force with the system in motion.

Because of its importance, the partial lubrication regime has been studied extensively. Efforts have been dedicated to experimentally characterizing this region, but in principle it would be preferred to have a way of parameterizing this curve without

extensive testing. Several, for example Gelinck et al. [17], have applied the known physics of the situation to develop models to predict this curve. Unfortunately these methods require a detailed knowledge of the surfaces in contact, including not only geometry, loading, and material properties but also a thorough characterization of the asperities, including average asperity dimensions and asperity density. Unfortunately this degree of understanding is not reasonably attainable in most machine design problems.

At velocities greater than the Stribeck velocity, full fluid lubrication dominates. This means that the system is hydrodynamically supported on a cushion of fluid. The system converges to a viscously damped system, with friction force increasing close to linearly with velocity.

Although it provides a good start to understanding machine friction, the dependence of force on constant velocity does not tell the whole story. The literature indicates that there are significant dynamics in the friction force response to changes in velocity. These dynamics relate to breaking away from stiction, as well as to changes in velocity while moving.

Dwell time has been shown to have an effect on the breakaway or stiction force. Within some finite time frame, ranging from seconds to hours, if two bodies stay in contact longer, it will take a larger force to produce slip. This effect, however, is not as significant as the rate of force application in breaking static friction. If the applied force is increased at a higher rate, less force will be required to induce slip [4]. Neither of these factors is generally of huge importance, but they are certainly present. They become especially interesting when stick-slip motion is examined in detail. Stick-slip motion is a result of the fact that a greater force is required to break away than to maintain motion.

One important effect in stick-slip behavior as well as any changes in velocity during motion is a time lag in changes in friction force. In a landmark paper [18], Hess and Soom report on an experiment in which they study the changes in friction force when velocity was varied rapidly. They find that the coefficient of friction (effectively, the friction force) might have a different value at the same velocity, depending on whether

the velocity was increasing or decreasing, and on the rate of increase or decrease. In other words, the coefficient of friction is dependent on not only the current velocity but also prior velocities. Hess and Soom show this effect to be attributable to a time lag between a change in velocity and the corresponding expected change in friction. They measure this time lag to be between 3 and 9 msec for their apparatus, depending on normal force and lubricant viscosity. Others have found similar values for this delay [4].

Integrated Friction Models

The ever-improving collective understanding of friction has led many to develop models for friction through the years. The strongest models integrate each of the effects described above to provide a single expression that, if properly parameterized for a particular apparatus, can predict friction under any conditions. Such models build on the partial models of earlier pioneers such as Dahl, Hess and Soom, and frequently compare to their experimental results.

Armstrong presents a model that incorporates most of the characteristics described above but that is forced to switch between what are essentially two different models, one for sliding and one for pre-sliding behavior [3]. Others argue that there is no physical basis for such a switching approach, that qualitative analysis of friction clearly shows that friction is a continuum of intertwined effects [25].

Canudas de Wit et al. present a formulation that does not require a switching function between the pre-sliding and sliding regimes [6]. This model is widely acknowledged in the field and is sometimes referred to as the LuGre model. The model can be easily reduced to the Dahl model, and the authors demonstrate through simulation that their model can show the time lag behavior described by Hess and Soom, as well as a dependence of static friction force on the rate of force application. The model does not show a dependence on dwell time. It does, however, exhibit stick-slip behavior that is qualitatively similar to observed physical results. As stated above, the model properly exhibits the dynamics and the Dahl effect. What is of most interest here, however, is the steady state force-velocity relationship. This model predicts

a steady state force F_{ss} given by the following relationship:

$$F_{ss} = F_c \text{sgn}(v) + (F_s - F_c) \times e^{-\left(\frac{v}{v_s}\right)^2 \text{sgn}(v)} + bv \quad (5.9)$$

In equation 5.9, F_c is the Coulomb force, F_s is the static friction force, v_s is the Stribeck velocity, and b is the damping coefficient. This curve fairly effectively predicts the behavior shown in Figure 5-3.

Swevers et al. offer an improvement to the LuGre model [25]. They argue that the LuGre model is insufficient in the pre-sliding regime, as it offers fewer degrees of freedom than the number of actual distinct physical phenomena that influence friction. They point out that the LuGre model dictates the shape of the force/displacement curve in the pre-sliding regime, and it therefore cannot be made to match experimental data. Swevers et al. introduce an improved model that allows for hysteretic behavior in the pre-sliding regime and the opportunity to fit this curve to specific experimental data. Since the sliding regime is of much greater concern here, this distinction is largely academic for the purposes of this analysis.

Model Simulation

Along with models for friction comes a desire to simulate these models. While the integrated models described above show proper behavior in both the friction/velocity and friction/position relationships, they are discontinuous at zero velocity. This can prove very problematic for simulation. Many have suggested, among them Bauchau [5], the use of some continuous approximation of the functions for very low velocities. This might involve simply a linear region with a very high slope connecting the breakaway forces for positive and negative velocities. Thus zero velocity is actually only permitted at zero force, but the velocity will remain very low unless force exceeds the stiction level. Bauchau addresses the effect of selecting the characteristic velocity at which a transition is made from the linear region to the Coulomb region. He points out that the selection of this velocity is essentially a tradeoff between smoothness for computational ease and model validity.

Parameter Identification

Using friction models to describe behavior in actual devices leads to a need for identifying the model parameters. Good models are presented with such experimental data to validate themselves, but unfortunately this is not found in every paper. A number of people have attempted to identify certain parts of friction in actual hardware; only a handful have reported truly thorough friction characterization. If viscous damping is largely the only form of friction and the system is linear, parameters can be identified with classical methods such as frequency response or step response. The difficulty comes in identifying nonlinear systems, a field that encompasses most devices with substantial friction.

The describing function approach has been used in order to allow the use of frequency response analysis for nonlinear systems [4, 20]. The work by Kang et al. [20] is one example of this approach. Kang et al. apply this approach to identify the friction in a precision positioning system. The goal is to predict and compensate for friction to aid in a position feedback control system. The authors focus largely on identifying the stiction regime and the Dahl effect, and use a simple Coulomb model in the sliding regime.

Some of the most thorough friction identification is reported by Armstrong [3]. Armstrong details a series of experiments on a PUMA robot. He identifies the friction/velocity curve, examines position dependence of friction (which in his case is not substantial), performs stick-slip testing, and observes the expected dynamics of changing velocity, varying dwell time, and applying force at varying rates. Armstrong uses position, velocity, and force measurements to obtain estimates of acceleration, allowing him to identify the portion of the applied torque that contributes to acceleration, to load bearing, and to overcoming friction.

Armstrong finds that friction is extremely repeatable in his robot when he considers both position and velocity dependence. This repeatability is shown even after a period of months and thousands of operating hours. Position dependence in the PUMA amounts to only a small correction on the velocity dependent friction force.

Armstrong's findings are, to the author's knowledge, the most complete and carefully documented example of machine friction characterization available in the literature.

5.2.2 Model Simplifications

The exceedingly complex world of machine friction can and must be simplified for the purposes of the current application through careful consideration of the needs and characteristics of the design. It should be noted that the rotary-to-linear transmission used in the module offers a significantly higher degree of complexity than the systems typically studied in the literature. For example, Hess and Soom studied a single line-contact interface between a rotating disk and a stationary part. The carefully instrumented apparatus provided detailed information about the motion and forces in each body. Their hardware is typical of that used in the tribology community. Even Armstrong's PUMA used only a simple spur gear train, and this robot was manufactured with precision tolerances.

The module studied here has many moving parts, most of which are not instrumented at all. The rollers in the nut, the ball contact and wiper contact in the linear bearing, the bearings that guide the screw, are all moving and interacting, and no position or force information is available. Additionally there is deliberate compliance in the screw mounting, and known backlash in the Rollnut. All that is known is the rotational position of the screw and the endpoint force. Thus a friction analysis requires a series of significant assumptions.

Much of the friction identification literature focuses on studying stiction, the pre-sliding machine, and stick-slip behavior. Because of the heavy position dependence of friction in this device, any study of stiction would require examining the conditions at each point along the screw length, and the motion in each direction. A conventional stick-slip experiment would be nearly impossible, as the flow of the experiment might be dictated almost exclusively by position dependent friction. Before embarking on such a difficult endeavor, some thought should be given to the use of the robot. As opposed to most devices described in the literature, this robot is not designed to provide precise position control. Most position controlled robots, including

Armstrong's PUMA, are always driven from one side. Using knowledge of the pre-sliding regime to help break static friction is a good idea, because it is obvious when movement is desired - this is dictated by the controller. (See the following chapter for a more detailed discussion of friction compensation techniques.) In the case of an impedance controlled device, it is the interaction between robot and environment that is regulated, not position. Thus it may be desirable to break static friction in either direction at a given time, depending on the interaction forces. This subtlety certainly complicates the idea of using a model to help break static friction, and calls into question the feasibility of doing so effectively. In light of the difficulty of creating such an experiment, and considering the feasibility questions, it was decided to forego any experiments dedicated to the pre-sliding regime. Perhaps a reasonable estimate of breakaway force can be projected from the results of friction/velocity testing. It should be noted as well that since the pre-sliding regime is ignored, the effects of dwell time and rate of force application are also ignored.

The focus of the experimentation on friction described here is on understanding friction force as a function of position and velocity. Because of the strong position dependence, it is not sufficient to generate a curve analogous to Figure 5-3. Curves like this must be generated along the entire length of the screw.

Perhaps the most significant assumption that must be made regarding the determination of the friction/velocity and friction/position relationships is that the endpoint friction can be treated as a single, lumped friction mode. When the nut translates up the screw, there are a number of points of contact between multiple bodies, moving at different speeds. For example, the rollers must cover more distance riding around the screw thread than the balls moving directly vertically up the linear bearing. Each of these interfaces contributes simultaneously to the endpoint friction force, and each in principle has an associated Stribeck curve like that in Figure 5-3. The Stribeck velocity and the shape of the curve are likely different for each mode, especially when considered as a function of endpoint velocity. It is clear from Gelinck's analysis [17], for example, that the Stribeck curve is also dependent on loading, which will vary throughout the device and can change as the robot is loaded. It is assumed, however,

that the sum of all modes acts as a single friction mode; essentially the movement of the nut is viewed as the motion of a block sliding up or down against a surface. All friction modes are considered lumped into the single available set of data.

5.2.3 Experiment Design

An experiment was devised to characterize the friction in the screw module as a function of position and velocity. As mentioned in the previous section, the system was examined from the point of view of the handle. The friction at a given position and velocity is considered as a single force value representing a lumped combination of the friction from each component, including reflected friction torque from rotating components as well as friction force in the translating components.

Testing was performed on the complete assembly, including the following potential sources of friction: the nut/screw interface, the bearings in the nut, the bearings guiding the screw, the linear bearing, and the motor. The testing was performed with the module in a passive mode, with the amplifier on and enabled, but zero torque commanded to the actuator. Note that the motor “friction” torque can include components of actual mechanical friction as well as additional damping due to the motor resistance.

The experiment was designed to drive the system at constant velocity from the endpoint. The force to maintain this constant velocity motion was measured with the ATI force transducer. Using constant velocity ensured that none of the input force was used to accelerate the system mass. Because $\ddot{z} = 0$ and $F_c = 0$, equation 5.3 reduces to

$$F_f(z, \dot{z}) + mg = F_p \quad (5.10)$$

By driving the system in several different directions and configurations, F_f and mg can be determined.

The system was driven using a Parker positioning stage. The stage was driven with a Compumotor stepper motor. The motor actuated 25000 steps per revolution, and the stage advanced 5 mm per revolution. A picture of the experimental setup




Figure 5-4: Experimental setup for driving the system at constant velocity.

is shown in Figure 5-4. A piece of angle iron was attached to the positioning stage. Force was transmitted to the nut through a mounting piece attached to the force transducer. A closer view of this interface is shown in Figure 5-5. Driving downward required crossing zero force input. Thus it was not sufficient to simply “push” the system in one direction to apply the force. For this testing, a coupling was included between the transducer and the angle iron on the stage. This coupling included two ball joints to allow the transmission of forces in the desired direction but prevent forces in other directions from building up due to slight misalignment between the driving and driven systems. This coupling is shown in Figure 5-6.

Testing was performed across a broad range of speeds, from 0.5 mm/sec to 150




Figure 5-5: Interface for force transmission for friction testing against gravity.




Figure 5-6: Double ball-joint interface for force transmission for friction testing with gravity.

mm/sec. 0.5 mm/sec is far slower than a person will typically move for any sustained period of time. 150 mm/sec was the maximum speed at which the stage would successfully actuate. This value was expected to be far beyond the Stribeck velocity, and should thus suffice to project the viscous damping coefficient to estimate friction at higher velocities, as the operating range will exceed 150 mm/sec.

Testing Upside Down

To determine the force due to gravity in the system, some testing was repeated with the system turned upside down. The gravitational force on the system, mg in Equation 5.3, is independent of position and velocity. Thus the force to drive the system at a given position and velocity is equal to the force to overcome friction plus mg . With the system turned upside down, the direction of gravitational force is reversed, with respect to the motion of the system. So the force then equals the friction force minus mg . Thus the difference between the force measured under these two conditions is $2mg$.

Assuming there is no direction dependence on friction in the screw system, this data could be determined from simply looking at the data from driving the system upward and downward, without turning it upside down. This simple test provides an extra check on the direction-dependence of friction, and should increase confidence in the determined gravitational force.

5.2.4 Experimental Results

Frequency Domain Analysis

The constant velocity data provides a great deal of insight into the dependence of the friction on position as well as velocity. Friction is often studied as a function of velocity; indeed this idea is central to the design of the friction experiment. Because of the strong position dependence in the system studied here, however, it is important to consider the velocity dependence locally with respect to position. In other words, the behavior shown in Figure 5-3 is expected, but is understood to be the variation in

friction force as the nut passes a particular point on the screw at different velocities.

Some effort is required to understand how tightly discretized the position must be considered to be. To make a first estimate, it is important to understand the expected sources of friction, based on an understanding of the geometry and on qualitative feel from the prototype module. Clearly there is a component of the force, due to misalignment of the screw and bearing, that is constant, or varies close to monotonically with position. Any bowing of the screw could introduce friction that varies with the frequency of the screw's rotation. This friction mode would also include friction in the bearings in which the screw rotates, as well as the motor bearings. Each of these modes can readily be felt in the screw module.

The friction force data was examined in the frequency domain to observe the effects of the two modes described above and to look for additional modes. A Fourier transform of the force data was taken. Figure 5-7 shows the force data in the frequency domain, with a best-fit trendline removed to reduce the extreme low frequency end of the plot. Figure 5-7 shows the data taken at 20 mm/sec, and is representative of the data at all velocities in most essential features. Note that there is a substantial amount of noise to at least several hundred Hertz. This is likely sensor noise, as the ATI force transducer is known to produce noisy data, particularly at such a high resolution. The spike at about 200 Hz is not common to all data and is likely not traceable to any real effect.

Figure 5-8 shows the same data as Figure 5-7, but is focused only on the lower end of the frequency spectrum in order to point out several interesting details. Note that there is a substantial spike just above 1 Hz. Recall that the nut travels 19.05 mm per revolution, and that this data represents a test performed at 20 mm/sec. From this it is clear that the spike at 1.05 Hz represents the mode at the frequency of the screw's rotation.

There is another, smaller spike visible in Figure 5-8. This spike occurs at about 4.2 Hz, roughly 4 times the frequency of rotation of the screw. Recall from chapter 3 that the effective diameter of the screw is approximately 14.2875 mm (0.5625 in), and that the effective roller diameter is approximately 3.5 mm (0.1378 in). Since the




Figure 5-7: Friction force magnitude versus frequency, traveling at 20 mm/sec.




Figure 5-8: Friction force magnitude versus frequency, traveling at 20 mm/sec.

roller diameter is roughly one fourth that of the screw, this means that the rollers rotate four times for every revolution of the screw, or at four times the frequency of the screw. This indicates that the spike at 4.2 Hz in Figure 5-8 is due to friction in the bearings that guide the rollers. This assumption is supported by the fact that there is no spike at twice the screw's rotational frequency, so the spike at 4.2 Hz appears to be a real phenomenon, not just a harmonic of the screw's vibration.

The frequency domain analysis revealed the expected modes at low frequency and at the frequency of the screw's rotation, as well as an unexpected but understood mode at the frequency of the roller rotation. There were no other significant frequency components found to be common to all velocities. Thus the significant friction modes are present and understood.

Data Processing

With some understanding of the friction modes that are desired to be preserved, as well as of the actual frequency content of the data, steps can be taken to make sure that the desired information is extracted from the data. It is certainly desirable to quantify the low frequency friction properties, including the magnitude of friction variation with rotation of the screw. It could prove desirable to also understand the effect of variation in the roller friction. From the frequency domain analysis, however, it is clear that this behavior is not nearly as significant.

It is unrealistic and somewhat unproductive to consider substantially higher frequency content than that mentioned here. The small magnitude of these changes means it is difficult to distinguish them from measurement noise, and any change in the module's friction properties over time may significantly alter these small variations. Furthermore, measurement and computation limitations make it rather unlikely that such high frequency effects could be realistically compensated for.

With measurements at 14 different velocities and along the entire length of the screw, there is a significant amount of data to be handled, and including more high frequency data greatly increases complexity. The force data was low-pass filtered at a frequency one decade above the frequency of rotation of the screw. Thus the impor-




Figure 5-9: Optimal low pass filter used on force data taken at 25 mm/sec.

tant frequency components are preserved, and unwanted noise is removed. Filtering was performed off-line using Remez optimal filtering techniques. For each velocity, a 400 point filter was designed permitting a 30 Hz “don’t care” band after the cutoff frequency. A characteristic filter is represented in the frequency domain in Figure 5-9. This is the filter used for data at 25 mm/sec. The targeted cutoff frequency is at about 13 Hz.

To look at velocity dependence for each point, it is necessary to define regions in position for which friction is assumed to be constant. Any points sampled within each region or “bin” can be averaged to obtain a force measurement for the bin. Initially the screw was split into 142 bins from bottom to top, with each bin 2.5 mm long. The rationale for selecting this bin size follows loosely from Nyquist’s sampling theorem: It is desired to detect frequency components up to 4 times the rotation frequency. To detect this, one needs to sample at twice this rate, or approximately 8 times per revolution of the screw. Recall that one screw revolution covers 19.05 mm, so each revolution contains almost 8 bins 2.5 mm long.

Some experimenting was done with other bin sizes, and the data appeared best and easiest to understand with 2.5 mm bins. This bin size is used throughout the




Figure 5-10: Friction force versus velocity, 25 to 50 mm from bottom of screw.

analysis.

Driving Upward

The force data for driving the screw upward was processed as described in the previous section. Note that the test was performed against gravity, so that the gravity force is added to the actual friction force. The results for gravity identification will be discussed shortly, but it is useful to know that the gravitational force is around 10 N.

Figures 5-10 and 5-11 show two representative cases of friction force versus velocity for the upward travel. Each figure shows the curves for 10 adjacent bins spanning 25 mm of the screw's length. Figure 5-10 covers the range from 25 mm to 50 mm, while Figure 5-11 covers the range from 250 mm to 275 mm. These figures demonstrate a number of interesting points about the friction in the module and about the experiment used to gather the data.

The curves in Figures 5-10 and 5-11 are laid on top of each other not with the intent of confusing the reader, but rather to show the general characteristics of each of several regions of the screw. Due to the erratic nature of even the processed data,




Figure 5-11: Friction force versus velocity, 250 to 275 mm from bottom of screw.

curves for individual bins can be misleading. By viewing several curves together one can attempt to understand the data in terms of its overarching trends.

Figure 5-10 exhibits the expected characteristics discussed in the friction background section. The friction force appears to approach some static level at low velocities. At low velocities friction force appears to decrease slightly with velocity, exhibiting negative viscous damping behavior. The Stribeck velocity appears to be at approximately 20 mm/sec, and beyond that value the friction force increases with velocity. One might strike a line through the middle of the data and suggest that the force increases by roughly 5 N from 25 mm/sec to 150 mm/sec, indicating a viscous damping coefficient of 40 N/m/sec.

Figure 5-11 certainly shows a different characteristic. At velocities beyond 20 mm/sec the behavior is similar, but the force has greater magnitude; the force increases with velocity, perhaps even with a similar slope. This is not surprising; one might expect the increased force not to result from viscous damping, but from a sliding component. There is no more grease in the upper half of the screw, but the normal forces are higher and Coulomb friction becomes greater.

At lower speeds, the friction force continues to decrease as velocity decreases,

and in fact appears to do so more rapidly. This is unexpected behavior, but the trend appears strongly whenever measurements are made on the upper half of the screw. This behavior is not obviously explained by friction theory. It is important to remember, however, that the foundation of Figure 5-3 is in the idea of a single friction mode, of one body moving against another. The system tested here is orders of magnitude greater in complexity, and contains a number of interfaces all moving at once. Perhaps each friction mode independently behaves as predicted, but each has a dramatically different Stribeck velocity and hugely different parameters. One could imagine that, with the limitations of the experiment done here, the understanding of each mode reflected to the endpoint results in the data shown in Figure 5-11. Beyond this uncertainty, the data is not fully understood.

Less heavily discussed in tribology literature but perhaps more significant here is the dependence of friction force on position. Figure 5-12 shows the friction force versus position at 50 mm/sec. Several trends are quite evident from this plot. Momentarily ignoring the oscillations and considering the gross force trend, there is a distinct ramp increase in friction force between 100 and 200 mm. Beyond 200 mm, the force stays fairly constant. This interpretation matches qualitative evaluations of the friction force; it is much easier to move the handle when the nut is on the lower half of the screw.

The most regular oscillations in Figure 5-12, those with the greatest amplitude, are at the frequency of the screw's rotation. The screw can rotate between 18 and 19 times in the total travel of the nut. The amplitude of these oscillations seems to roughly increase with the overall amplitude of the force.

The region between 0 and 50 mm in Figure 5-12 shows another increase in friction force. This increase is not seen in the first set of data taken. A second set of data was taken at velocities from 40 to 150 mm/sec in an attempt to obtain better data at higher velocities. This increase was seen in the new data, from which Figure 5-12 is taken. This change is a bit mysterious, and is discussed more later.

Figure 5-13 shows all the force versus position curves from the first set of data for velocities from 0.5 to 50 mm/sec. With a few exceptions, the data points fall in a




Figure 5-12: Friction force versus position at 50 mm/sec.

fairly distinct band of data. This band represents the variation in force with velocity. The relatively small width of the band shows that friction is, for this device, much more dependent on position than on velocity, at least for smaller velocities when viscous damping does not play a large role. Figure 5-13 shows that the force is fairly repeatable with position, across several velocities, and presents hope that the force can be predicted closely based on position. Note that since Figure 5-13 is taken from the original set of data, the increase in force in the bottom 50 mm of the screw is not present.

High Velocity Testing

In all cases, the individual curves do not seem to offer good data at the higher velocities. This held true throughout the testing. While an explanation for this is not immediately apparent, one artifact of the experiment does suggest that data at lower velocities might be better. Recall that each data point represents an average of all the force measurements across a given length of position at a certain velocity. Since the sampling rate is always the same, when the system is moving at lower velocities




Figure 5-13: Friction force versus position, velocities from 0.5 to 50 mm/sec.

it records more data points in each bin. Thus there is a larger sample size making up each data point. Since the data is filtered well below the Nyquist frequency in all cases, it is difficult to understand why this should be a problem, but it is one possibility.

Despite the problems with testing at the higher velocities, the data at lower velocities appears to be, for the most part, smooth and continuous. To illustrate the smoothness at lower velocities and relative choppiness at higher velocities, Figure 5-14 shows a surface plot of friction force versus both position and velocity. Note that the left side of the plot, where velocity is less than 50 mm/sec, shows significant variation of force but does so in the absence of any large discontinuities. The right side, on the other hand, shows some extreme irregularities, and substantial discontinuities with respect to both position and velocity. If Figure 5-14 is an accurate depiction of the friction in a real system for all velocities observed, a reasonably smooth surface would be expected.




Figure 5-14: Friction force versus position and velocity, upward travel.

Repeatability

Because of the huge amount of data and the time required for each test, an extensive analysis of the repeatability of friction was not feasible in the time frame of this analysis. Ideally one might measure friction a number of times to increase statistical confidence in the data. The scope of this project instead is limited to understanding the basic characteristics of the friction in the device. Repeatability is, however, an issue, and should at least be mentioned in passing.

In his analysis of a PUMA robot [3], Armstrong-Helouvy finds that friction is quite repeatable, even after thousands of hours of use and many months. This is, of course, dependent on consistency in all other variables (some of which, such as dwell time, are ignored here).

As mentioned in the previous section, some of the data was taken twice in an effort to improve friction measurement at higher velocities. In these cases, the measurements for the two sets of tests were compared in order to determine the repeatability. One such comparison is illustrated for 50 mm/sec testing in Figure 5-15. Note that where the data looks good, the two curves tend to match up. This is particularly true




Figure 5-15: Friction force versus position at 50 mm/sec, two separate trials.

between 100 and 300 mm. As mentioned before, the second curve shows an increase in friction at the extreme low end of the screw. This suggests a possible slight change in the device, but otherwise the measurements look fairly similar.

This brief sidelight does not attempt to represent a thorough examination of repeatability of friction. It instead serves only to suggest that where tests were repeated, results looked similar. Unless something occurs to suggest the contrary, it is likely acceptable to conditionally assume that the test data represents an accurate depiction of the characteristics of the device, conditions that are likely to be repeated if necessary.

Driving Downward

The data for driving the nut downward showed some of the same characteristics as for driving upward, with a few important exceptions. Since the geometry is not, by design, direction dependent, the force to push the nut down is expected to be the same as the force to push it up, except in the opposite direction and corrected for the effects of gravity. In effect, this means that if a horizontal line is struck in




Figure 5-16: Friction force magnitude versus velocity, 25 to 50 mm from bottom of screw, downward travel.

Figure 5-12 or in Figure 5-10 at the value of the gravitational force (roughly 10 N), the corresponding curve for downward motion is expected to be a reflection of the original curve about that line.

Figures 5-16 and 5-17 show the friction force versus velocity going downward for the position ranges corresponding to Figures 5-10 and 5-11. Note that Figure 5-16 shows something close to the expected trend, as there is a force spike at low velocities. The Stribeck velocity is a bit difficult to discern due to the deviation from the curve between 10 and 25 mm/sec. This strange shape in the curve appears to a lesser degree in other position regions as well. Again, this phenomenon is not immediately explicable, but may result from the superposition of several friction modes.

Figure 5-17 shows reasonably close correlation to Figure 5-11. Again, there is an unexplained drop in force at very low velocities.

More interesting than the differences between the friction versus velocity plots with change in direction are the differences between the friction versus position plots. The friction force versus position at a velocity of 50 mm/sec in the downward direction is shown in Figure 5-18. Here, since the movement is downward, following the plot




Figure 5-17: Friction force magnitude versus velocity, 250 to 275 mm from bottom of screw, downward travel.

in time actually means moving from right to left on the plot. Note that again, the oscillation with rotation of the screw is present and significant, and that the force profile at the lower end of the screw is similar.

There is a significant difference, however, in the gross force versus position trend in Figure 5-18 versus that seen in Figure 5-12. While traveling upward, there is a distinct ramp increase in force until about 200 mm, where the force seems to level off, with the exception of the oscillation. Traveling downward, however, the corresponding force increases more slowly through the middle of the screw but continues to ramp up all the way to the top of the screw. This point is perhaps better illustrated by presenting both curves together. Figure 5-19 shows the original curve for upward travel as a solid line, and the opposite of the curve for downward travel as a dashed line, shifted by about 20 N to roughly compensate for twice the gravitational force. The difference in the curves is subtle, but it is present in data taken at each velocity. This suggests that the friction is, in fact, direction dependent.




Figure 5-18: Friction force versus position at 50 mm/sec downward.




Figure 5-19: Friction force versus position at 50 mm/sec upward (solid) and at 50 mm/sec downward shifted by 20 N (dashed) to match the upward curve.

Upside Down

To further investigate, and to assist in identifying the gravitational force in the system, some testing was performed with the module turned upside down. This reverses the direction of the gravitational force with respect to the screw. For simplicity of fixturing, and because a full set of data is not needed to answer the questions at hand, testing was first only performed for a segment of the screw from about 190 mm to 325 mm. This section coincides with the region where the two directions appear most different in their friction characteristics. Note that the position is still defined the same way with respect to the screw, and moving in the positive direction still means moving the nut toward the motor. So, with the module upside down, moving in the positive direction means moving physically downward. The data presented here involves moving each way at 5 mm/sec, and is representative of the other speeds tested as well. The data was gathered, processed, and binned as described before.

Figure 5-20 offers a comparison of the force to move the module at 5 mm/sec in the positive direction, right-side up and upside down. 20 N (roughly twice the gravitational force) has been added to the curve taken from the upside down data to show the two curves on top of each other. Note that the upside down data, shown by the dashed line, increases more with position than does the right-side up data.

Figure 5-21 compares the force magnitude to move the module at 5 mm/sec in the negative direction. A similar shift has been applied to the data from right-side up measurement. In this case, the data taken right-side up, the solid line, shows a more dramatic force magnitude increase with position.

Figure 5-22 compares the curve showing movement in the negative direction right-side up with the curve showing movement in the positive direction upside down. Note that in both of these cases movement is with gravity. The two curves match more closely than the curves in the previous two figures. A similar result is observed in comparing the two sets of data for moving with the assistance of gravity.

This difference, primarily visible qualitatively in these plots, is supported statistically. The standard deviation of the difference between the two curves in Figure




Figure 5-20: Friction force versus position at 5 mm/sec in the plus direction, with module right-side up (solid) and upside down (dashed). The upside down curve is shifted by 20 N for comparison.




Figure 5-21: Friction force magnitude versus position at 5 mm/sec in the minus direction, with module right-side up (solid) and upside down (dashed). The right-side up curve is shifted by 20 N for comparison.




Figure 5-22: Friction force magnitude versus position at 5 mm/sec with gravity, with module right-side up (solid) and upside down (dashed).

5-22 is only 1.1 N. For Figures 5-20 and 5-21 the standard deviation of the difference between the two curves is 1.74 and 1.33 N, respectively. This data provides an extremely interesting result. The friction force appears to be more heavily dependent on the direction of gravity with respect to the motion than on the direction with respect to the screw.

There is at least one plausible explanation for this behavior. The force is applied to the robot through a force transducer mounted with its center at a distance of 77.8 mm (3.065 in) from the center of the screw. Recall that any force applied vertically at the transducer produces not only a reaction force but also a reaction torque at the screw and bearing assembly. This is illustrated in Figure 3-3 and discussed in Chapter 3. The linear bearing is intended to absorb these parasitic torques and to primarily transmit only vertical forces to the Rollnut, producing relative motion against the screw.

It is evident from looking at any of the data provided in this chapter that the forces to move against gravity are substantially greater than those to move with gravity; in fact, the force to move with gravity is generally very close to zero, because

the gravitational force is very similar to the friction force. With a lower force at the transducer, the parasitic torque on the bearing and nut is far lower. Perhaps some of the torque is indeed transmitted through the bearing to the nut, and the different loading condition results in different friction characteristics. This could produce results like those seen here.

In principle, one could design an experiment to study the correlation between applied parasitic forces and torques and changes in friction force. The complexity of such an experiment, however, would be nearly prohibitive. The experiments here deal with only a vertical force applied at the endpoint. There is no reason that forces and torques in all six axes could not be applied at the endpoint, and the effect of each studied thoroughly. This level of depth is not necessary here. The essential components of friction are apparent and consistent. Any compensation method that involves a model will be an approximation at best, and slight differences will not be a severe problem.

Gravitational Force Identification

Because the friction cannot quite be identified independent of gravitational force, the planned method for more precisely measuring the effect of gravity is not as effective as hoped. It is important to note, however, that the forces are considerably lower in some portions of the screw than in the region examined in the previous section. If the theory expressed in that section is correct, at a low enough force the differences due to parasitic torques will vanish. The simplest way to identify the gravity term is then to repeat testing upside down in the lower friction region of the screw, and to compare this data to the prior data moving against gravity.

The experiments with the module upside down were repeated, this time with travel focused on the lower half of the screw. Data was taken working with and against gravity in both directions of motion. Results are considerably better than those taken along the upper half of the screw in that the force/position curves are very similar in shape, and are simply shifted by twice the gravitational force. The mean of the difference between the two curves taken across all trials, for motion in




Figure 5-23: Friction force magnitude versus position at 5 mm/sec in the minus direction, with module right-side up and upside down. The curve against gravity (dashed), is shifted by $(2 \times mg)$, or 20.075 N.

both the positive and negative direction, was found to be 20.075 N, with a standard deviation of 1.056 N. Thus the gravitational force mg in equation 5.3 was found to be half of this difference, or 10.04 N. This is consistent with the sum of the masses of translating parts being approximately 1 kg.

Figure 5-23 shows the force data for driving the screw in the negative direction with and against gravity. The data taken against gravity, the dashed line, is shifted by 20.075 N.

5.2.5 Model for Friction in the Screw Module

The preceding sections have been devoted to identifying F_f and mg in equation 5.3. In order to perform experiments to identify M , and to simulate the system, it is necessary to model these identified terms.

It has already become clear from the experimental data that equation 5.3 is not in fact a full and accurate description of the system. The friction force, represented by F_f , is apparently a function not only of z and \dot{z} but also of the loading conditions.

Since it is difficult and probably not worthwhile to identify the effect of these loading conditions, the basic format of equation 5.3 will be maintained. Because the data, however, represents the best picture of what is going on in the system, it is best to use all the data gathered in creating a model. In spite of the fact that data taken driving the system upward may not be an entirely accurate depiction of friction *only*, it is true that when the robot is used, at least without active control, those same loading conditions will be present. That is to say, if an upward force is applied to the handle, there will be a force and a torque induced at the screw. So the force data for pushing the nut in both directions will be used in modeling the system. Since this is the case, there is no point in attempting to decouple the separately identified mg parameter from the friction force. This parameter is included as part of F_f for this portion of the analysis.

The friction force must be fitted as a function of position and velocity. It is preferable for several reasons to fit with a function that is decoupled in position and velocity if a good fit can be found. This reduces the computational difficulty of the curve fitting, and also helps to retain some physical meaning in the model. Both position and velocity dependence are understood, to some degree. Functions are selected that take advantage of the physical intuition associated with the characteristics of the friction force.

Model for Upward Motion

The data from experiments where the nut was driven upward, against gravity, can be fitted to a function of the following form:

$$F_{fu} = F_{vu}(\dot{z}) \times F_{pu}(z) \quad (5.11)$$

The position dependence is broken into two functions:

$$F_{pu} = F_{pu1} + F_{pu2} \quad (5.12)$$

F_{pu1} provides the gross trend in friction force with position along the screw. F_{pu2} provides a sinusoidal component that mimics the oscillation in friction force as the screw revolves. The velocity dependent function F_{vu} serves as a scaling factor, correcting the more significant position dependent term for changes in velocity.

F_{pu1} was obtained using a seventh order polynomial fit to the position data averaged over all velocities. The data and the fitted curve are shown in Figure 5-24. The frequency for F_{pu2} follows fairly obviously from the geometry, and the phase angle for the sinusoid can be quickly determined by examining the data. The amplitude of this sinusoid as a function of position was fitted iteratively with a fifth order polynomial. The resulting function is shown, again with the averaged force data, in Figure 5-25. Note that while the sinusoid fits well in frequency and phase at certain parts of the curve, at other points it does not match the experimental data as well. This indicates that the frictional peaks due to binding of the screw as it rotates do not always occur at the same angular position as the nut translates up the screw. This is not entirely surprising, but the approximation is fairly accurate here. The model used best matches the data at the large oscillations near the center of the workspace, where most of the activity will be.

After an acceptable fit for friction versus position at an average velocity was determined, a function to adjust for different velocities was included. No attempt was made to model the effects of stiction or any local maximum in friction force at low velocity. Because such a pattern was not seen along the entire screw, and because dealing with friction at such low velocities leads to a flurry of problems in measurement and simulation, a model was adopted that results in zero friction force at zero velocity. To make a reasonable fit under these conditions, 10 N was subtracted from the friction force to remove the effects of gravity. This step is necessary because only the friction should be scaled with velocity, not the gravitational force. The 10 N for gravity is added back into the model after the position term is scaled. By some simple graphical analysis it was decided that the F_p curve should be scaled by about a factor of 1.3 at 150 mm/sec, and by roughly a factor of 0.75 at the lowest velocities tested. A piecewise curve was fit to the friction/velocity data, averaged




Figure 5-24: Force versus position, averaged across all upward velocities, with a polynomial fit (dashed) of the gross position trend.




Figure 5-25: Force versus position, averaged across all upward velocities, with a fit (dashed) of the gross trend and the variation with screw rotation.

over position. The curve and the fit are shown in Figure 5-26. For velocities greater than 2.5 mm/sec, a line with a slope of 0.00296 N/mm/sec was used. This gives a good idea of the viscous damping coefficient. For lower velocities, a power curve was selected to create a continuous function that goes through zero, and to slightly smooth the discontinuity in the first derivative where the two curves meet. The actual fits for each of the curves discussed here are given below. z and \dot{z} are used here in mm/sec. Note that the basic equation has been slightly modified from 5.11 to reflect the 10 N shift for gravity.

$$F_{pu1} = -7.27 \times 10^{-16} z^7 - 9.97 \times 10^{-14} z^6 + 6.63 \times 10^{-10} z^5 - 3.54 \times 10^{-7} z^4 \\ + 7.57 \times 10^{-5} z^3 - 0.0069 z^2 + 0.2334 z + 13.6820$$

$$F_{pu2} = (1.17 \times 10^{-11} z^5 - 8.14 \times 10^{-9} z^4 + 1.42 \times 10^{-6} z^3 + 2.13 \times 10^{-6} z^2 \\ + 5.55 \times 10^{-7} z - 0.0052) \sin(\omega z + \phi)$$

$$F_{vu} = 0.754 \times \dot{z}^{0.15} \text{ for } \dot{z} \leq 2.5$$

$$F_{vu} = 0.00296 \dot{z} + 0.857 \text{ for } \dot{z} > 2.5$$

$$\omega = \frac{2\pi}{19.05}$$

$$\phi = 0.385$$

$$F_f = F_{vu}(F_{pu} - 10) + 10$$

Figure 5-27 shows plots at two different velocities of the processed friction force minus the modeled friction force versus position. These plots are fairly typical of all velocities. Note that the fits are far from perfect; this is because both in position and velocity, fits were made to average curves that cannot capture all the details of each different curve. In most cases, however, the error is between -2 and +2 N. Perhaps a better fit could be obtained with a coupled function of position and velocity, but this is a satisfactory level of error for this task. The fit is not good at the extreme ends of the screw, but the problems are for the most part confined within the extreme several centimeters, and should not present too great a problem for the time being.




Figure 5-26: Force versus velocity, averaged across all position bins, with a piecewise fit.




Figure 5-27: (Friction force)-(modeled force) vs position at 5 mm/sec (left) and at 20 mm/sec (right).




Figure 5-28: Force versus position, averaged across all downward velocities, with a fit (dashed) of the gross trend and the variation with screw rotation.

Model for Downward Motion

The same procedure was followed to fit the data taken from driving the nut downward. Again the position curve at average velocity was fitted with a sinusoid on top of a seventh order polynomial. One interesting problem did arise in the course of the position fitting. As in the prior data set, the sinusoid does not always match up with the data. The problem is a bit more extreme in this case, however, and some of the peaks are as much as half of a revolution off the typical trend. This is most likely again explained by differences in the loading conditions. After some experimentation, a 180 degree phase shift was included in the sinusoid at all positions below about 212 mm. There is no real physical basis for this, but the adjustment seems to fit the data fairly well. The only part of the fit that is much weaker occurs very close to the phase shift, from around 190 mm to 230 mm. The friction/position data averaged across velocity and the associated fit are illustrated in Figure 5-28.

The velocity fit was implemented in the same way as before, again with the phase shift. The friction/velocity data averaged across position and the curve fit are shown

in Figure 5-29. The full set of equations quantifying the fit of the data driving downward is as follows.

$$\begin{aligned}
F_{pd1} &= 4.76 \times 10^{-15} z^7 - 5.56 \times 10^{-12} z^6 + 2.50 \times 10^{-9} z^5 - 5.39 \times 10^{-7} z^4 \\
&\quad + 5.66 \times 10^{-5} z^3 - 0.0028 z^2 + 0.0676 z + 2.7001 \\
F_{pd2} &= (2.59 \times 10^{-11} z^5 - 1.49 \times 10^{-8} z^4 + 2.01 \times 10^{-6} z^3 + 2.59 \times 10^{-6} z^2 \\
&\quad + 1.14 \times 10^{-7} z - 0.0038) \sin(\omega z + \phi + \pi) \text{ for } z \leq 212.5 \\
F_{pd2} &= (2.59 \times 10^{-11} z^5 - 1.49 \times 10^{-8} z^4 + 2.01 \times 10^{-6} z^3 + 2.59 \times 10^{-6} z^2 \\
&\quad + 1.14 \times 10^{-7} z - 0.0038) \sin(\omega z + \phi) \text{ for } z > 212.5 \\
F_{vd} &= 0.8098 \times \dot{z}^{0.15} \text{ for } \dot{z} \leq 2.5 \\
F_{vd} &= 0.0025 \dot{z} + 0.923 \text{ for } \dot{z} > 2.5 \\
\omega &= \frac{2\pi}{19.05} \\
\phi &= 0.385 \\
F_f &= F_{vd}(F_{pd} - 10) + 10
\end{aligned}$$

Several representative plots of friction data minus the model are shown in Figure 5-30. The quality of the fit is comparable to that found in driving upward.

5.3 Inertia Identification

Given the dominant nonlinearities in the system of equation 5.3, identifying the effective endpoint mass is a difficult task. With a simple linear system a broad set of tools in both the frequency and time domains are available for identification. For this system, a simple $F = ma$ approach was taken in the time domain.




Figure 5-29: Force versus velocity, averaged across all position bins, with a piecewise fit.




Figure 5-30: (Friction force)-(modeled force) vs position at 5 mm/sec (left) and at 40 mm/sec (right), downward velocities.

Component	Endpoint Mass (kg)
Bearing carriage	0.3
Nut	0.29
Mounting Bracket	0.18
Force Transducer	0.05
Motor rotor	2.13
Screw	3.25
Coupling	0.22
Total	6.42

Table 5.1: Estimated contributions to endpoint mass, by component.

5.3.1 Estimation

The known properties of the components can be used to obtain a reasonable estimate of the endpoint mass. This mass consists of the mass of parts that translate as well as the reflected mass of parts that rotate.

Parts that translate upstream of the force transducer include the carriage of the linear bearing, the nut, and the mounting bracket. Part of the force transducer mass should also be included. Rotating parts include the screw, the coupling, and the motor rotor. Less significant parts, such as the inner race of the bearings and the rollers in the nut, are considered negligible.

The mass of the carriage is provided in the product specifications. The nut and mounting bracket were weighed, as was the force transducer. The force transducer outboard mass was determined by Doeringer [12] and from this the inboard mass is determined. The rotor inertia is known and the coupling and screw inertias can be estimated based on their geometry and material densities. The reflected mass can then be easily computed with equation 5.4. The mass contributions of each component are presented in table 5.1. The total estimated reflected mass is 6.42 kg. This includes the mass of everything on the actuator side of the force transducer. The patient also has to move the mass of the handle assembly, but this mass cannot be decoupled from the mass of the patient hand and arm from a modeling point of view.

5.3.2 Experiment

The system was back-driven by hand in generally oscillatory motions at a variety of speeds and amplitudes. The endpoint force and the position were recorded as always.

Extracting the Desired Parameter

To determine the mass parameter M from equation 5.3 it is necessary to identify each of the other terms. A nonzero acceleration is required to apply Newton's law. As long as acceleration and all other terms can be determined, data from any motion of the module can be used to compute the mass.

In order to reduce the number of variables and therefore the uncertainty, F_c was set to zero for these experiments. Force F_p was applied at the handle and measured with the transducer. Velocity and acceleration were determined from the time history of the position measurement. The friction force F_f was determined from the model described above.

The basic idea of the experiment was to use equation 5.3 to estimate the $M\ddot{z}$ term at each moment in time and to divide by the computed acceleration \ddot{z} to evaluate the mass.

Velocity and Acceleration

The determination of velocity and, in particular, acceleration from discretely sampled position data is non-trivial. It is nearly impossible to avoid some phase lag in finite-difference differentiation. This phase lag is exacerbated by the application of low-pass filters that become necessary due to quantization of position data that effectively quantizes velocity and acceleration [12].

For this experiment much of the difficulty was avoided because the data was processed offline, making past, present, and future data points available. At each point k a first approximation of velocity was determined using a center-difference calculation:

$$\dot{z}_k = \frac{z_{k+1} - z_{k-1}}{2T_s} \quad (5.13)$$




Figure 5-31: Magnitude versus frequency for the low-pass filter used to process velocity and acceleration.

This estimate was refined by applying a 400 point optimal filter, designed with the `remez` command in Matlab. The cutoff frequency used was 30 Hz, with a “don’t care” region extending to 50 Hz. The magnitude frequency response of the filter is shown in Figure 5-31.

The acceleration was then computed in a similar fashion, by center-differencing the velocity and running again through the same filter.

Despite the use of the center-difference method and careful filtering, it is difficult to place too much confidence in data that is twice-differentiated from discretely sampled data, particularly at low values of the derivatives. Furthermore, since the acceleration is the divisor in the mass calculation, at extremely low accelerations, errors can lead to instability. Because of this concern, samples where the acceleration dropped below 2.5 mm/sec^2 were not used for mass identification.

Application of the Friction Model

The position- and velocity-dependent model for friction and gravitational force in the module, as formulated in the preceding sections, was applied to the position and

velocity data to calculate F_f . Recall the several assumptions made in the creation of this model. Although all the data used to create the friction model was taken at constant velocities, it is assumed that the same friction force will extend to the accelerating system, as it passes through each velocity. Thus the dynamics of friction are neglected.

The model makes no attempt to show stiction behavior and is therefore inaccurate at low velocities, with the error reaching a maximum at zero velocity. Using the model below the Stribeck velocity would introduce substantial additional error, so samples where the velocity magnitude was below 20 mm/sec were omitted.

Figures 5-27 and 5-30 document some of the error between the friction model and the sampled data to which it was fitted. This error, generally less than 2 N in magnitude, extends to the application of the model to equation 5.3. Thus the possibility of an error of at least ± 2 N is expected in the $M\ddot{z}$ term. This error could be even larger at velocities beyond those tested in the friction identification phase. The mass data cannot be determined from any one sample or small set of samples, but must instead be determined through examination of a large number of samples to improve statistical certainty.

5.3.3 Experimental Results

Following the procedure documented in the preceding paragraphs, the time history of $M\ddot{z}$ and \ddot{z} were generated for each data file. Ideally, these two curves should simply be scaled versions of one another. Recall, however, that there is some uncertainty in each of these curves. Figure 5-32 shows a representative example of these two curves from one of the test data files. When the oscillations are rapid and large in magnitude, the curves show very similar characteristics.

Figure 5-33 offers a closer look at the same data, with the $M\ddot{z}$ curve in this case scaled by a factor of 0.2 to offer a better visual comparison. It is obvious from this plot that the acceleration leads the $M\ddot{z}$ curve by a fairly consistent amount, here approximately 8 msec. This relationship is fairly typical for these data files, and is likely an artifact of the processing of the data.




Figure 5-32: Mass times acceleration (solid) and acceleration (dashed) versus time.




Figure 5-33: (Mass times acceleration)/5 (solid) and acceleration (dashed) versus time.

Data Processing and Statistical Analysis

Not surprisingly, trials with high velocities and accelerations, generally with a rapid frequency of oscillation, tended to produce the most usable data. For each suitable data set, the optimal time shift was found to make the two curves line up. This was done by trying a broad range of time shifts, and selecting the shift that minimized the standard deviation of the ratio of the two curves. 5 data sets were considered. For each set the best shift was found to be between 16 and 36 samples, or between 8 and 18 msec.

No physical meaning can reasonably be attached to the need to align the curves in this fashion. This process does seem, however, to in some ways compensate for some of the modeling errors that might result from the processing of the acceleration data or the estimation of friction.

Once the curves were properly aligned, the unusable points with low velocities and accelerations were neglected and the curves were divided. Figure 5-34 shows a typical sampling of the inertia data versus time. The acceleration curve is plotted along with the data. The points to be neglected show an inertia of zero in this plot.

The values of inertia across all 5 data sets were averaged. This average of course only included data points where the acceleration and velocity were both above the minimum required value. The mean of all points in all sets is 5.49 kg, with a standard deviation of 1.72 kg. This is about 15% less than the estimated value. The discrepancy could be due to errors in estimation or in friction model inaccuracies. The tests to determine inertia used velocities substantially greater in magnitude than 150 mm/sec. If the viscous damping coefficient projected in the model is too high, this would result in friction force estimates that are too great and an estimate of endpoint mass that is too low. Other errors in the friction model could have a similar impact on the determination of endpoint mass. Clearly the accuracy of the inertia identification is closely tied to the accuracy of the friction identification.




Figure 5-34: Inertia and acceleration versus time. If velocity or acceleration is too low, inertia is shown as zero.

5.4 Subsystem Conclusions

Understanding the transmission of the module is heavily dependent on gaining an understanding of its friction characteristics. While enough data has not been gathered to fit all the dynamic parameters of a modern friction model, the relevant dependencies on position and velocity are satisfactorily understood. The model proved good enough to predict the friction force to within a reasonable error of the measured value, and to obtain at least a good ballpark estimate of effective endpoint mass. This model can now be applied to the controller development process. If further detail is needed about the system, additional tests can be performed at a later date.

Chapter 6

Controls Analysis

The screw module was designed to provide mechanical properties such that, in the presence of an impedance controller, it would provide stable interaction spanning a broad range of endpoint impedances. To achieve a satisfactorily low impedance, it may be useful to incorporate advanced control techniques. This chapter addresses several potential control methodologies for reaching the desired performance. Having characterized the system, it is now useful to study several different control architectures to consider their viability in achieving the desired control of this particular system.

6.1 Linear Impedance Controller

The MANUS robots have used a simple linear PD controller to great success in regulating endpoint impedance, despite the nonlinearities in the system [15]. It should be noted that the PD type of impedance controller makes no attempt to increase or decrease the effective inertia of the physical system, but for MANUS this value is an acceptable one. Attempting the use of a similar controller is a sensible starting point for controller design on the new module.

A PD controller of the form:

$$F_c = -K_c(z - z_d) - B_c(\dot{z} - \dot{z}_d) \quad (6.1)$$

where K_c and B_c are the controller gains, was applied to the system. Briefly assuming that the system is linear and applying the definition of impedance for a linear system given in Chapter 3, the endpoint impedance can be expressed as:

$$\frac{F_p(s)}{V(s)} = Ms + (b + B_c) + (k + K_c)/s \quad (6.2)$$

where M , k , and b are the mechanical system's effective endpoint mass, stiffness, and damping, respectively (Gravity is momentarily neglected for simplicity). Since the system has no preferred position, there is no stiffness so $k = 0$. The mass is simply the mass identified in the previous chapter, and the apparent endpoint damping represents a sum of the derivative controller gain and the system damping. Of course b in actuality represents a complex nonlinear function of position and velocity so that, as the system is nonlinear, it is in fact not appropriate to use the Laplace domain. Equation 6.2 is simply used to illustrate the fact that the total system impedance is a combination of the existing mechanical impedance and the software impedance, and cannot realistically be reduced below the value dictated by the mechanical parameters.

The system was tested with a PD controller. In addition to the PD controller, a constant force of 10 N was applied to compensate for the bias due to gravity. The robot will primarily be used either at zero impedance or with a fairly soft spring and possibly some damping. In clinical trials with MANUS, 100 N/m is a typical spring constant. Note that all the testing in the prior chapter, with zero control effort, is identical to testing the module with a PD controller with zero gains. So all the data in Chapter 5 gives a good idea of zero-impedance performance.

Figure 6-1 shows a force versus position curve for the module with a 100 N/m spring centered at 0.180 m commanded in software, and zero damping. The desired curve is also presented. This data was taken by actuating the device slowly by hand, and no attempt has been made to adjust the curve by considering the portion of the input force used to accelerate the mass. This acceleration is visible in the curved portions of the data near the endpoints.

While the desired slope is roughly present, it is clear from Figure 6-1 that there is




Figure 6-1: Force versus position, actual (solid) and desired (heavy dot). PD controller. Arrows show direction of travel, circle denotes starting point.

a large, undesired force offset roughly 10 N in magnitude depending on the direction of motion. This is the effect of the substantial friction force. This is confirmed upon viewing another plot, this time at a higher commanded stiffness. Figure 6-2 shows the results of a similar test with a stiffness of 1000 N/m commanded. Again, the offset from the desired curve is roughly 10 N in each direction - the approximate magnitude of the friction force. (The large bulge in the curve at the extreme left side is due to actuator saturation.) In these and subsequent plots, the circle marks the starting point and the arrows mark the direction of travel in time.

One might say qualitatively that the performance shown in Figure 6-2 looks closer to the desired than that shown in Figure 6-1 (with the exception of the actuator saturation). This is supported by the “feel” of the device; in the case with the stiffer spring, the module has more the feel of an undamped spring, while with a soft spring the dominant feel is one of friction. Some of the difficulty in overcoming friction is lost to the fact that high forces are expected from the spring. This idea is reflected, in a way, in equation 6.2. If K_c is large, the percent change in impedance due to an increase in b is small. This trend continues with higher impedances, and




Figure 6-2: Force versus position, actual (solid) and desired (heavy dot). PD controller. The bulge at the lower end is due to actuator saturation.

confirms a concern from the design phase of this project: A transmission with a large amount of friction has little trouble offering a high endpoint impedance. The difficulty comes in emulating a low impedance. 100 N/m is within the desired operating range of the robot; 1000 N/m is likely not. This spurs further investigation into active compensation techniques to reduce the effective endpoint impedance.

Testing by Hand and Interaction Stability

The results in the preceding section, like all the experimental data presented in this chapter, were gathered from experiments where force was input to the module by hand. Hand testing can produce somewhat erratic results; it is difficult for a human to act effectively as either a force source or a velocity source. However it will be shown that hand actuation offers some advantages to the alternatives. To investigate impedance of a manipulator it is important to consider the impedance, or more correctly the admittance, properties of the environment.

Using instead the stage apparatus from the preceding chapter would mean that the robot, with an impedance controller, is interacting with an extremely stiff envi-

ronment. This touches on the issue of coupled stability. The case of an impedance-controlled system interacting with a stiff environment has been shown to in some cases lead to instability [8]. Coupled stability is an important and interesting issue, but it is not rigorously addressed here. Certainly the stability limits of any candidate controller should be understood, but initial evaluations must not necessarily include the worst case scenario.

In the absence of a thorough study on coupled stability, it is best to examine system behavior coupled to the closest available environment to its expected environment of operation. As the system is designed to interact with a human, the tests are conducted in this fashion.

Interaction with a human arm can be understood by examining simple system models. Figure 6-3 shows a bond graph of the system interacting with a human arm. This illustrates the use of a simple PD controller, and as such represents a model of the system used to gather the data in Figure 6-1. For this single degree of freedom system, a simple spring-mass damper model of the human arm is used, with parameters k_{arm} , b_{arm} , and M_{arm} (these parameters are in reality state-dependent). Note that the arm mass is lumped with the system endpoint mass, as they are assumed to move together; essentially the arm mass is assumed to be entirely located at the hand. The trajectory that the human is commanding is shown by the virtual trajectory \dot{z}_h , at the bottom of Figure 6-3. The desired trajectory \dot{z}_d dictated by the robot is shown at the left of the figure. Note the similarity in the structure of these two controllers; the mass is essentially subject to two superposed PD controllers. The system itself includes the total mass of the module and the arm, as well as the gravitational force and the friction force. To the left of the heavy dashed line is the impedance controller. Everything to the right essentially appears as an admittance to this controller.

A model of the desired system behavior is shown in Figure 6-4. The gravity and friction forces are gone, and the mass has been reduced to only the arm mass. The only difference between the force dictated by the impedance law and the force of interaction comes in accelerating the arm mass, which is in actuality part of the human system.

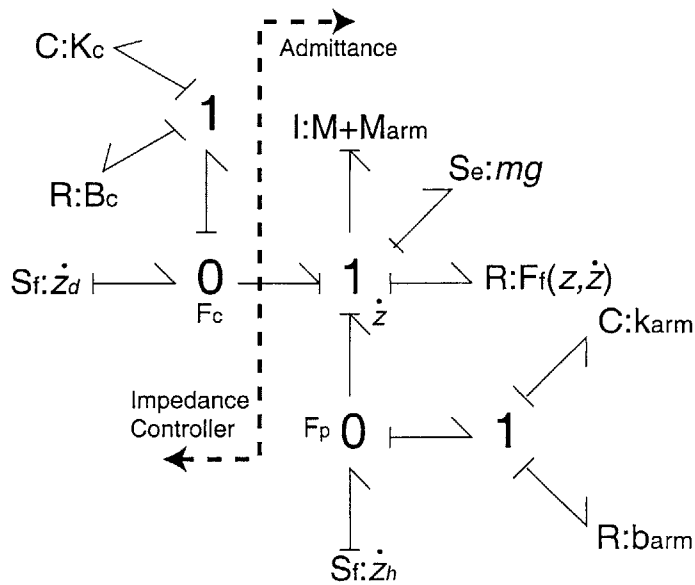


Figure 6-3: Bond graph model of the system under PD control, interacting with a human arm.




Figure 6-4: Bond graph model of desired system performance.

The deviation of the data in Figures 6-1 and 6-2 from the desired is a result of the friction force and the force needed to accelerate the mass (The effect of the gravitational force is removed). The arm dynamics are present as well but are downstream from the force transducer; the ideal manipulator system would provide exactly the correct relationship between position and force. When the device friction causes deviation from this relationship, however, the arm dynamics play a role. This must be kept in mind when examining the data presented here. The data should be viewed somewhat qualitatively and relatively; performance improves as the actual curve approaches the desired curve, for a given command stiffness.

6.2 Friction Compensation Techniques

As friction is such a common machine design problem, several approaches have been suggested to attenuate its negative effects. Not surprisingly, nearly all of these methods have been geared toward position- and tracking-controlled systems. Armstrong provides a comprehensive overview of all but the most recent developments in friction compensation [4]. A sampling and explanation of some of these techniques is provided here, loosely grouped by those which use strictly a feedforward approach and those which require some force feedback. A brief discussion of the applicability of analogous techniques to the impedance control of this specific device follows in the next section.

6.2.1 Feedforward Controllers

Stiffening the System

For a trajectory controlled system, many of the problems due to friction, particularly stick-slip, can sometimes be virtually eliminated by simply increasing the controller gains. Armstrong demonstrates that stick-slip behavior can be eliminated by implementing proportional and derivative gains that exceed some critical value, depending on the system parameters [3]. Clearly this amounts to increasing the impedance of a PD controller.

Dither and Impulsive Control

Dither and impulsive control are two techniques dedicated largely to reducing the effects of static friction. Dither is a continuous oscillatory force or torque command, added on top of the normal control signal, that keeps the system from ever coming to a complete stop, therefore eliminating problems from static friction. The magnitude of the dither force or torque must exceed the static friction levels.

Impulsive control is the use of a single pulse or a series of pulses to produce motion from a system at rest. The system is characterized and tested to determine the proper magnitude and duration of the impulse. Ideally, a large enough force is applied to make the system start moving, and then the controller transitions into another type of control that regulates the motion.

Model-based Control

The development of detailed models for friction based primarily on velocity has led to the use of feedforward friction compensators. The idea of such a compensator is to offset the instantaneous friction force by using the actuator to provide an equal and opposite force. Such compensators use friction models of varying complexity; some implementations use only a Coulomb friction model, others a Coulomb plus viscous model, and others a full-fledged, integrated friction model.

Armstrong, Canudas de Wit, and Swevers [3, 6, 25] each suggest a compensator of the latter style employing the friction models that they propose. In each formulation except Armstrong's, position dependence of friction is neglected for simplicity, so velocity is the only variable that must be fed back to estimate friction. Canudas de Wit mentions that his model requires information about a state that is not measurable (the average deflection of the individual asperities), so a state observer is incorporated to estimate this value [6].

The primary difficulty with any model that uses velocity to estimate friction force, occurs at zero velocity, where the friction force is undefined and can cover a large span. For trajectory control this is not a huge problem if the system is well modeled; the

desired direction of motion is known, so a force just larger than the static friction force must be applied to “jolt” the system out of the stiction regime.

6.2.2 Feedback Controllers

Torque or Force Control

If torque or force measurements are available, it is possible to use this information to enhance system performance in dealing with friction. Several researchers have engaged in direct joint torque control to reduce the effective friction in a geartrain, making a gear motor appear as a direct drive motor to the higher level position controller. The actual joint torque at the output of the geartrain is measured and compared to the desired joint torque dictated by the position controller, and additional actuator torque is added as needed. Reductions in friction torque of a factor of 30 have been reported [4].

The use of torque or force control in conjunction with a higher level task controller raises several control issues that must be examined. Generally the actuator and torque sensor are non-collocated; in fact this is a necessary condition of effective torque control. The torque control loop can only reduce the friction caused by elements between the actuator and the sensor. Additionally the use of two control loops requires multi-loop analysis unless sufficient frequency domain separation is maintained [4]. These and the other issues associated with force control in reducing the effects of friction are not well documented, and demand further study.

Adaptive Control

Another friction compensation method that requires sensor feedback is adaptive control. One of several adaptive algorithms can be used to continually adapt the model parameters in a friction estimator like those discussed in section 6.2.1. This adaptation can be done offline in an experiment or online throughout the operation of the machine. The structure of the controller is similar to the feedforward model-based approach, but the actual friction values change based on the adapting model

parameters.

6.3 Applying Friction Compensation to an Impedance Controller

The application of some of the compensation techniques used in position-controlled systems to impedance control is a simple conceptual step; for other techniques, careful attention must be paid to the differences between the two methodologies and to how the compensator affects the interaction port.

6.3.1 Impedance and Interaction

The fundamental difference between position control and impedance control is that while a position controller seeks to minimize a position error, an impedance controller seeks to minimize the error in the relationship between force and motion at the endpoint. Each type of controller may suggest a desired position or trajectory. In a position-controlled system, any deviation from this desired position, regardless of the interaction forces, is considered an error. In an impedance controlled system, if the deviation from the desired position is accompanied by the appropriate interaction force between the manipulator and the environment, there is no control error.

Stiffening a controller to minimize the tracking error from friction requires altering the impedance, which is in fact the measure to be controlled in an impedance-controlled system. It is undesirable for the friction compensation technique to require a certain impedance to work properly. It is potentially useful to know, however, that beyond a certain impedance stick-slip effects will not be a problem.

Before considering the other feedforward techniques suggested for trajectory control it is useful to consider again the nature of the therapy robot discussed here, to gain insight by considering its intended use. Examining again the system free-body diagram in Figure 5-1, it is important to understand that the motion of the mass is dictated not only by the control input F_c , but also by the input from the environment

F_p . That is, the system must be capable of being driven from both sides. The proper motion of the system is dictated by both forces in the sense that F_c might want to push the system one way based on the position and the impedance law, but F_p may overpower F_c and produce motion in the other direction. This scenario may very well describe desired system behavior - again, the idea of impedance control is that the force that the system gives to the environment is related properly to the motion that the environment provides.

Because the system must be driven from both sides, it is unreasonable to use techniques such as dither or impulsive control. To keep the system constantly in motion would be impractical, because the magnitude of the dither force would have to offset not only static friction forces but also the maximum potential environment force that might attempt to hold it still, and the impedance law would undoubtedly be violated. Furthermore, it is not known when the system should deviate from a rest position and in which direction, because this is dictated by not only the control law but also the environment force.

This final limitation also presents a flaw with the use of feedforward model-based compensation. Any such friction estimator that does not rely on interaction force information will fail at zero velocity in an impedance-controlled system. As previously noted, at zero velocity the actual friction force can vary from the breakaway force in the negative direction to the breakaway force in the positive direction. To properly follow the impedance law, it is impossible to know the correct time or direction to break out of the stiction regime without knowledge of the interaction force. Thus in practice any such compensator must be made continuous through zero velocity and will be ineffective in dealing with stiction. Despite this problem, the use of knowledge of the system is appealing as a means of reducing the effects of friction while the system is in motion.

6.3.2 Feedback Control

Since the idea of joint torque or force control is fundamentally geared toward correcting deviations in the force from that dictated by the control law, it can be applied

more directly to impedance control without suspicion. It is best that the force sensor be located as close as possible to the interaction port, so that it returns the actual interaction force. If the difference between the actual force and the force dictated by the impedance law goes to zero, the impedance error is, by definition, zero.

6.4 Compensation Applied to the Module

6.4.1 Negative Software Damping

The only way to attempt to improve low-impedance performance within the framework of a PD controller is to use negative values for the derivative gain, essentially adding negative viscous damping and reducing the middle term in equation 6.2. This technique will do nothing, even in theory, to reduce the effective mass, but could in theory substantially reduce viscous damping.

For stability it is essential that the term $(b + B_c)$ remain greater than zero for all possible system states. Since b in this case represents a complicated friction function, it is nontrivial to ensure this. Viscous damping is not necessarily the dominant source of friction in this device, but it is substantial; the model fit to data in the previous chapter suggests damping coefficients of approximately 24 to 27 N/m/sec. In theory, negative derivative gains of lesser magnitude than 24 N/m/sec should permit stability. Gains greater in magnitude may work at low velocities and may offset some of the friction from other sources, but introduce the risk of instability at higher velocities, when the negative viscous term in the controller exceeds the actual friction in the hardware. This could be an acceptable implementation, however, if the velocities at which the controller becomes unstable are safely outside the operating range of the module (The robot will not, for instance, ever need to move at tens of meters/sec).

Linear viscous damping has no effect at zero velocity, and its effect is relatively small at low speeds. Even in theory, this approach does nothing to combat stiction and will be less effective at low speeds. Additionally this type of compensation does not attenuate the undesired variations in friction with position. It simply results in




Figure 6-5: Force versus position, with negative controller damping of 50 N/m/sec (solid), and desired (heavy dot). Arrows show direction of motion in time, circle marks start of trial.

a force from the actuator assisting motion, proportional to the velocity.

The module was tested with negative derivative gains in the PD controller. The experiment performed was the same as in section 6.1, with a 100 N/m spring centered at 0.18 m commanded as the desired impedance. A test was first performed with a derivative gain of -25 N/m/sec, enough to approximately offset the viscous damping. At low velocity this barely improved performance at all. This is probably because the viscous damping is a fairly small component of the total friction, particularly at low speeds.

The test was repeated with a gain of -50 N/m/sec, in an effort to compensate for additional friction. The results are shown in Figure 6-5. As the assisting force is dependent only on velocity, the spring behavior is almost lost. Slight variations in velocity are quickly amplified by the assisting force. The use of negative gains in the PD loop is not an effective means of friction compensation in this device.




Figure 6-6: System under PD control with friction and gravity compensators.

6.4.2 Feedforward Friction Compensator

The model for machine friction developed in Chapter 5 was used to create a friction estimator that was subsequently tested on the robot module. The estimator computes a friction force estimate as a function of the measured position and velocity (determined from the time history of position), and adds an equal and opposite force to the control input.

This behavior can be partially accomplished, at least for nonzero velocities, as shown in Figure 6-6. A friction estimator and a gravity compensator have been added to the controller such that they ideally cancel the effects of the actual friction and gravity. The unwanted terms and the compensator terms are shown as dashed lines. If the compensator exactly cancels out the unwanted forces, the only difference between this model and Figure 6-4 is the extra system mass. This is the goal of feedforward control using a friction estimator. This implementation is similarly illustrated in block diagram form in Figure 6-7. The idea is for the term produced by the block marked “Friction Estimator”, a part of the controller, to exactly cancel the term from the block marked “Friction”, a part of the physical system.

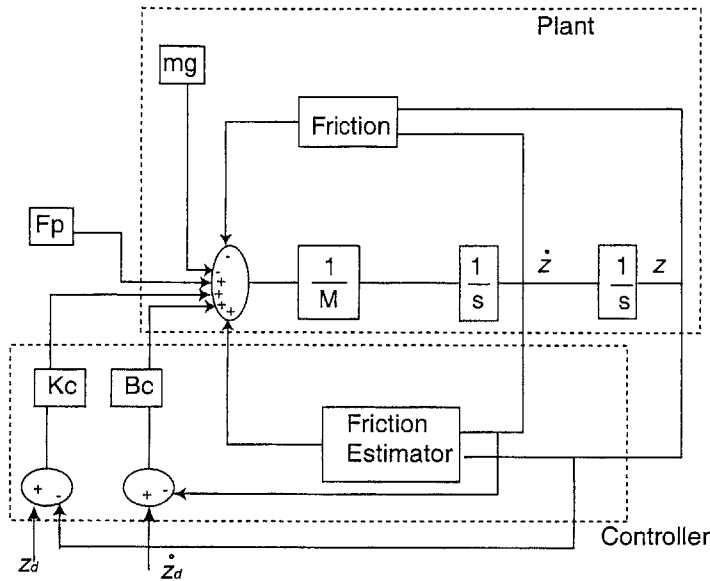


Figure 6-7: Block diagram of plant with a PD controller and a friction estimator.

The desired performance of an impedance controller is that the force exchanged at the port of interaction obeys the impedance law, in this case expressed by the gains of the PD loop. If the estimate is the exact opposite of the friction force, the sum of the forces on the robot inertia is simply equal to the PD input plus the interaction force. These two forces will only be equal and opposite if the mass is not accelerating. This approach does not offer any compensation for the nonzero system mass; hence the aforementioned difference between Figures 6-6 and 6-4.

Figure 6-8 shows the results of a test of the system with feedforward compensation, along with the prior results for the system with just a PD controller and the desired force/position curve. The test was performed at low velocity by hand, and the compensator provided a force equal to 90% of the friction force predicted by the model in the direction of motion. The performance appears to be better traveling down than traveling up; this is because additional mass in the form of a handle has been added that was not included in the gravity compensation, and this extra force effectively shifts the curve up. Clearly the compensator has improved performance. Because of the use of extremely low velocities, however, the assistance is limited. This is especially seen where the system stops. Note that at the endpoints of the plot, the




Figure 6-8: Force versus position, actual and desired (heavy dot). PD controller (solid) and feedforward friction compensator (dashed).

system reverses direction and passes through zero velocity. The distance between the two legs of the dashed line increases at the ends as the translation slows, reflecting the lack of compensation. This is particularly evident at the top end of the plot, where stiction must be broken to transition to downward motion. Figure 6-9 gives a closer look at this effect.

The rudimentary friction estimator used here is successful in using some knowledge of the system to reduce friction. Its effects are certainly more pronounced at high velocities. At high velocities this approach is largely limited only by the quality of the friction model.

6.4.3 Performance Enhancement with Force Feedback

One of the most appealing approaches to improving system performance is force feedback. This technique is particularly appealing because the control law does not require a system model. Keeping the PD controller to dictate desired impedance, an interior force control loop can be added to ensure that the force of interaction matches the desired force from the impedance law. This loop compares the interaction force




Figure 6-9: Closeup of Figure 6-8. Note the spike in force required to start downward motion, before the system reaches sufficient velocity for the compensator to take effect.

measured with the transducer to the desired force at any given time, and additional control effort is provided to compensate for the difference between the two.

Starting with the system equation of motion 5.3, the controller force F_c is now composed of two components, one part for the PD law and the other for force feedback:

$$F_c = F_{pd} + F_{fdbk} \quad (6.3)$$

where F_{fdbk} is the force feedback component. In equation 6.3 the symbol F_{pd} is used for simplicity, where

$$F_{pd} = -K_c(z - z_d) - B_c(\dot{z} - \dot{z}_d) \quad (6.4)$$

When the desired endpoint impedance is achieved,

$$F_p = -F_{pd} \quad (6.5)$$

Note that this is the true desired force of interaction, that could never be provided by a PD controller unless the system had zero mass or zero acceleration. This desired

condition is also demonstrated by the absence of the system mass M in Figure 6-4. With force feedback the effective mass can be reduced along with the friction. Achieving the equality in equation 6.5 suggests a form for the force feedback term in the controller. First applying a simple proportional controller, the control law for force feedback is:

$$F_{fdbk} = K_f(F_{pd} + F_p) \quad (6.6)$$

Plugging the controller defined by F_c into the equation of motion produces:

$$M\ddot{z} + F_f(z, \dot{z}) + mg = F_p + F_{pd} + K_f(F_{pd} + F_p) \quad (6.7)$$

This can be rearranged to:

$$\frac{M\ddot{z} + F_f(z, \dot{z}) + mg}{1 + K_f} = F_p + F_{pd} \quad (6.8)$$

It is evident from equation 6.8 that if $K_f \rightarrow \infty$, $F_p \rightarrow -F_{pd}$, and the desired impedance is achieved. Thus it is advantageous to use a high gain proportional control loop on the force to achieve desired performance, as long as stability can be guaranteed. In terms of the bond graph models presented previously, this control loop is applied between the forces at the two zero-junctions in Figure 6-3.

Equations 6.3 through 6.8, as they are presented, suggest that all that must be done to achieve perfect performance is to select a huge force gain. Much of import is left out here, and this is of course not the case. Force feedback is especially sensitive to sensor and actuator bandwidth and sampling rate. Simulation, experimentation, and practical insight show that the force gain must in reality be limited to moderate values.

A force feedback loop with the structure described above was implemented on the robot. Force data was filtered with a first order Butterworth filter with a cutoff frequency of 30 Hz. Data taken with a force gain of $K_f = 5$ is shown in Figure 6-10, along with the data from the unassisted PD controller. This controller clearly gives the best performance presented yet; the force offset is only perhaps 2 N on each side of




Figure 6-10: Force versus position. Force feedback controller with $K_f = 5$ (dashed) and PD controller (solid) with desired spring curve (heavy dot).

the desired curve. This makes sense, as the force gain of 5 should reduce the friction force by a factor of 5, from 10 N to 2 N.

The system feels easy to manipulate under this controller. The 2 N of unwanted force is nearly imperceptible, and stiction does not appear to be a problem.

The system was simulated in Matlab and Simulink to better understand the results of the force feedback testing. A block diagram of the system used for simulation is shown in Figure 6-11. Dashed boxes mark the plant, the human arm providing the interaction force, the PD controller, and the force feedback controller.

In the simulation, the PD controller used $K_c = 100$ and $B_c = 0$, just like the experiments. The system mass used was $M = 5.5$ kg. The human arm properties were $M_{arm} = 3$ kg, $k_{arm} = 300$ N/m, and $b_{arm} = 5$ N/m/sec. The virtual trajectory input by the human was a sinusoid 0.15 m in amplitude at a frequency of 0.8 rad/sec. The simulation included sensor resolution, sampling rate, and software delay consistent with those in the physical system. This simulation incorporated a model for friction based on the model formed in the previous chapter. The simulated model was modified slightly to include an approximation of the Stribeck effect and stiction.

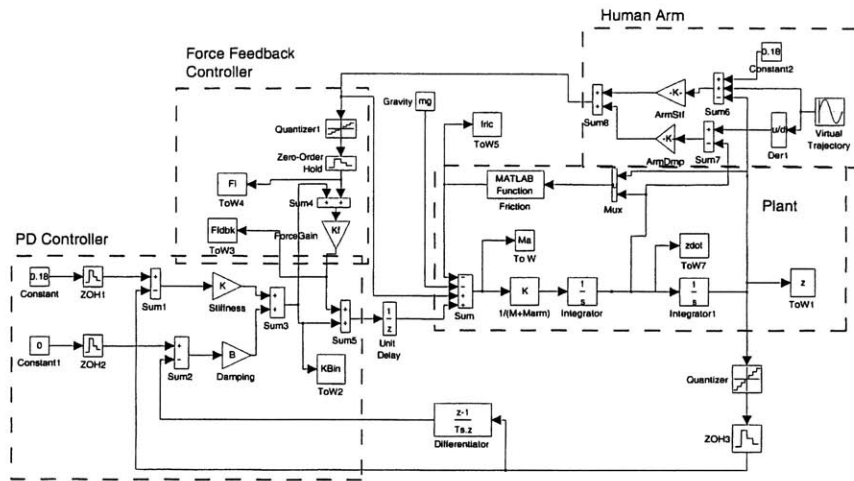


Figure 6-11: System simulated with force feedback. Plant, human arm input, PD controller, and force feedback controller are marked by boxes.

The system was first simulated with no force feedback, that is $K_f = 0$. The results of this simulation are presented in Figure 6-12. This curve shows many similar characteristics to the experimental data presented in Figure 6-1, including a hint of oscillations in friction force at a similar frequency. This is the variation in friction force as the screw turns. The experimental data includes more high frequency content. This is expected because it is known from the friction data in Chapter 5 that the actual friction has higher frequency components than does the model used in the simulation, and because the actual system is known to be of higher order than the model. This first simulation suggests that the model reflects system behavior reasonably well.

The system was then simulated with a force feedback gain $K_f = 5$, the same gain used experimentally to gather the results in Figure 6-10. The simulation results are presented in Figure 6-13. Note that in both cases stiction is more prominent in simulation than in experimentation. It is likely that the stiction value used in simulation is too high. Still, the simulation largely exhibits the same behavior observed in the physical system. Again, a reduction in the deviation from the desired curve from about 10 N to about 2 N is seen.




Figure 6-12: Force versus position, simulated with PD controller (solid) and desired (heavy dot).




Figure 6-13: Force versus position, simulated with PD controller with force feedback gain of 5.

Limiting factors in Applying Force Feedback

As previously described, higher force gains lead to better performance. It is desirable to implement the highest possible force gain that will give stable system performance with sufficient bandwidth. It is useful to consider the factors that limit force gains in real systems.

Actuator bandwidth, sensor resolution, sampling rate, and computation delay all affect controller bandwidth and stability. This system performs well with modest force feedback gains, but this may be largely because it uses extremely high resolution sensors and a high sampling rate (2 kHz). Additionally, the force feedback loop is not isolated. It is operating in the presence of a PD control loop, and the dynamics of the two loops may interfere with each other. Also present is the dynamic system of the environment, in this case the human arm, which may affect stability.

With force gains much greater than 5, chattering becomes a problem in the system. This chattering may be a result of the system backlash, as it is a more significant problem where the system is “looser” and the friction is lower. The simulation, which does not include a model for backlash, suggests that gains as high as 50 might be used in a force feedback loop without stability problems. This suggests that it is indeed an unmodeled effect, such as backlash or higher order dynamics, that is limiting the feedback gain.

If backlash is in fact shown to be the limiting factor for force feedback control, it may be advantageous to consider a preloaded nut in future versions. A preload will result in increased friction, but if its use permits a force feedback gain large enough to offset the additional losses, it might be worthwhile.

A more complete understanding of the factors limiting force feedback control would be useful in optimizing this design. Perhaps it is wasteful, for example, to have such a high resolution position sensor, because the sampling rate limits performance long before the limit due to position measurement is reached. A number of potential sources of limitation can be identified, but the importance of each source is not well understood as of yet.

6.4.4 Controller for the Screw Module

For the specific hardware discussed here, force feedback control offers solid performance, but model-based control methods also offer substantial improvements without relying on a high-gain feedback loop. It is tantalizing to consider a control structure that combines the two approaches. While the performance of the feedforward technique is insufficient in its own right, particularly at low velocity, perhaps this knowledge of the system can be applied to assist the feedback controller and improve performance. A feedback loop with a gain of 5 reduces force error by roughly a factor of 5; if the error is smaller to begin with because the friction has been predicted by an estimator, the final error will be reduced.

Perhaps the force information can be used in ways other than a proportional controller. PID or lead/lag methods might be investigated, or perhaps an observer can be created that uses the friction model as well as the interaction force to determine and compensate for the friction force in the stiction regime as well as the sliding regime.

With such a broad range of potential control architectures to choose from, the next step is to systematically analyze these ideas to determine the potential of each. This should be done to eventually arrive at the best available controller for this system, or at least to guarantee stability in the desired operating range at the desired performance.

Chapter 7

Conclusions

This thesis describes the development of a robotic module for upper limb therapy for motor recovery. The module, when mounted to an existing robot, will expand the planar capabilities of the current robot into a three-dimensional workspace, permitting the training of useful shoulder and elbow motions against gravity for recovering stroke patients.

A central focus of the module's design was maintaining the possibility of low endpoint impedance. This primarily meant minimizing friction and inertia in all three degrees of freedom. The screw-driven design maintained reasonably low impedance in all three directions of motion. The need for precise position alignment was avoided by introducing compliance into the mounting of the screw. For future versions of the design, such compliance should be considered before the parts are manufactured, and the use of compliance should be optimized. This may include placing the compliance differently or in more parts of the mounting. It is acceptable to trade position accuracy for lower endpoint impedance, to some degree.

The hardware and software of the system were carefully analyzed and characterized. The dynamics of the electronic system and the discrete nature of the software control are well understood, at least within the frequency range of interest. The current system sampling rate of 2 kHz approaches the maximum achievable rate with current computing power. A dramatic increase in sampling rate would most likely require a significant step up not only in PC power but in the speed of the digital

electronics of the force transducer.

The system transmission was carefully examined, particularly its friction properties. Friction was found to be rather difficult to measure, and experiments took a large amount of time. It would be preferable to repeat these experiments a number of times to improve statistical confidence in the results, but there was not sufficient time to do this. The friction data showed some of the expected characteristics, but not others. Friction was found to be largely position dependent as well as velocity dependent. Perhaps it would be better, from a control point of view, to have a system with more consistent friction, even if this friction is generally of greater magnitude.

Several methods of compensating for friction and inertia with active control were considered and attempted. The best results came from a simple proportional force feedback controller, though a feedforward friction estimator showed some promise. An expanded analysis of each of these controllers, and some consideration of the possibility of combining the two would be very interesting.

7.1 A Framework for Further Investigation

Study of this system has raised a number of closely related questions that demand further investigation. These questions primarily center around understanding the relationship between the mechanical impedance dictated by the uncontrolled hardware and the available range of impedances once active control is applied, and generalize to the design of virtually all interactive machines.

When designing an interactive machine one must define the desired interactive behavior of the device, including the range of endpoint forces and impedances, the range of motion, and the bandwidth. The machine is then designed to meet these requirements. It is likely that these requirements can be met through many combinations of hardware design and control strategy. One method is to design the hardware so that the device meets the essential requirements without control assistance. This could result in excessive cost in terms of mechanical complexity, or loss of performance in less essential areas. Another method is to design hardware that does not

meet the requirements, and to use whatever sensors and processors are required to bring performance into the desired range with force-based control methods, based on experience or trial and error. It is the author's opinion that this approach is followed by most designers of interactive machines and haptic devices, in the absence of any systematic method of determining the effects of control in altering apparent endpoint impedance.

This process could be greatly optimized with the presence of analysis tools to understand the relative importance of each factor in impedance-enhancing control methods. A designer of an interactive machine might then plan from the beginning on a certain hardware performance and a certain performance enhancement from software, and know approximately what sensors, actuators, and processors he will need to achieve this performance enhancement. This knowledge could deeply impact the hardware design itself, as the influence of properties such as the type and magnitude of friction, inertia, stiffness, backlash, etc. are considered. Perhaps it is far easier to apply force feedback to a system dominated by viscous damping than one with substantial Coulomb friction, even if the magnitude of the damping is much greater - then one might design for more lubrication and viscous contacts. How useful is it to fully identify the system? How much is gained from accurately modeling the system versus treating it as a "black box"? If the answers to these questions suggest that knowledge of the system is the most essential thing, design might focus on a machine that is easily characterized.

The field of interactive machinery could benefit greatly from a deeper understanding of these and related issues. Their study through experimentation, simulation, and theoretical analysis could yield great returns useful not only to rehabilitation robotics but to all those who study interaction of machines with an environment.

Ideally the use of active control to lower impedance should have been considered from early in the design phase of this project, as it may have permitted advantageous changes in the hardware design. This was not a realistic possibility, however, because of a lack of analysis tools for understanding the capabilities and limitations of force feedback. The compilation and verification of such a set of tools would be useful to

those designing and implementing interactive machines.

Appendix A

Compliant Mounting Analysis

Compliance was introduced to the screw mounting to reduce binding problems. The use of compliance is acceptable because position requirements are not extremely stringent. With compliance, however, comes the risk of dynamic resonance. This analysis considers this risk and methods for preventing resonance problems.

A.1 Model for Compliant Mounting

Potentially introducing compliance in both the top and bottom mounts produces the situation depicted in Figure A-1. The nut, which translates in z , acts as a moving pivot about which the screw can rotate. At the top and bottom are springs and dampers, marked k_1 and b_1 at the top and k_2 and b_2 at the bottom. These represent the compliant members that come between the screw and the frame, here shown as ground. At the top of the screw is M_{top} , representing the mass of any additional components that may be fixed to the top of the screw. The distance from the origin to the center of mass of this top mass is l_{cm} . The distance from the origin to the bearings is, in each direction, l_1 . The idea of such a configuration is that the nut determines the position of one point on the screw, and the two ends settle into the position of least resistance, minimizing binding.

The dynamics of the system can be determined by examining the equivalent inertia, stiffness, and damping about the pivot point as a function of z . Neglecting

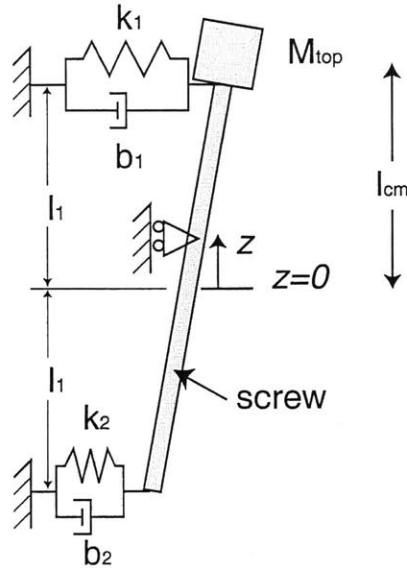


Figure A-1: Model for the system with compliance introduced in top and bottom bearing mounts.

damping for now, these can be shown to be:

$$I_{eq}(z) = I_{cm_{screw}} + M_{screw}z^2 + M_{top}(l_{cm} - z)^2 \quad (\text{A.1})$$

$$k_{eq}(z) = k_1(l_1 - z)^2 + k_2(l_1 + z)^2 \quad (\text{A.2})$$

where I_{eq} and k_{eq} are the equivalent rotational inertia and stiffness, respectively.

A.1.1 Driving Frequencies

It is important to consider what range of frequencies might be excited in the everyday use of the screw module. The system can then be designed to keep all resonant frequencies due to mounting compliance out of the operating range.

The system can be expected to see speeds up to about 1 m/sec. As the nut travels 0.01905 m per revolution, this amounts to a maximum rotational frequency of about 50 Hz. It is wise to keep all resonant frequencies at least a decade above this value, at 500 Hz or greater.




Figure A-2: Moment of inertia about pivot vs. nut location, Case 1.

A.2 Case 1: Motor and mount piece move with screw

One possibility is to place compliance between the mounting flange that holds the upper bearing bore and the frame of the robot. The motor is rigidly mounted to the mounting flange. The motor and flange are connected rigidly to the screw, so this is essentially the situation depicted in Figure A-1 with $M_{top} = M_{motor} + M_{flange} = 1.60$ kg and $l_{cm} = 0.344$ m. The moment of inertia of the screw about its center of mass is $I_{cm_{screw}} = 0.011125$ kg \times m² in the plane shown in Figure A-1.

This configuration results in a moment of inertia, according to equation A.1, that varies with nut (pivot) location as shown in Figure A-2. The inertia is greatest when the nut is at the bottom of the screw, because the mass concentrated at M_{top} is farther away from the pivot.

To maintain a minimum natural frequency of at least 500 Hz regardless of where the nut is along the screw while using the most compliance possible, one might use a greater stiffness at the top and a lesser stiffness at the bottom. Using, for example, a




Figure A-3: Natural frequency vs. nut location, Case 1.

stiffness $k_1 = 4 \times 10^7$ N/m at the top and $k_2 = 5 \times 10^6$ N/m at the bottom produces natural frequency as a function of nut position as shown in Figure A-3.

To use this design without placing a resonant frequency in the robot's operating range, it is necessary to use compliant members at least as stiff as the values suggested in the preceding paragraph.

A.3 Case 2: Motor and mount piece grounded

Another design for a compliant mounting involves mounting the motor and its mounting flange rigidly to the robot frame, and introducing compliance between the bearing that guides the screw and the frame of the robot. This produces relative motion between the motor shaft and the screw shaft; this motion is accommodated by the flexible coupling. This implementation means that in Figure A-1, $M_{top} = 0$ and the motor and flange mass become part of the ground.

Since $M_{top} = 0$, the inertia is only dependent on the screw inertia and screw mass and is symmetric along the top and bottom halves of the screw. Equal stiffnesses can be used at the top and bottom. For example, with $k_1 = k_2 = 2 \times 10^6$ N/m, the




Figure A-4: Natural frequency vs. nut location, Case 2.

resulting natural frequency as a function of nut position is shown in Figure A-4. The natural frequency stays in an acceptable range, using stiffnesses that are substantially lower than those required for case 1. Thus a “looser” system can be used, without lowering the natural frequency.

A.3.1 Damping

The possibility of adding damping was explored in an effort to flatten the curve shown in Figure A-4. Damping generally lowered the resonant frequency across all positions without lowering it substantially more at the center of the screw. Damping is disadvantageous because it increases friction in the system and increases the required stiffness to keep the natural frequency high enough.

A.4 Case 3: Tuned Vibration Absorber

A final possibility considered was the use of multiple compliances and masses to increase the order of the system. The motivation behind this was to use two resonances




Figure A-5: Model for system with tuned vibration absorber.

to cancel each other out, in other words to implement a tuned mass-damper system [24]. In such a configuration, an intermediate mass may be oscillating substantially, but the point of interest, in this case the handle, is isolated from such oscillation.

For this hardware a tuned mass-damper might be implemented by using compliance in both of the locations mentioned in the preceding sections. This configuration is shown in Figure A-5. M_{top} is again the mass of the motor and mounting flange, and k_1, b_1, k_2 and b_2 are placed as described for Case 2. k_3 and b_3 are added between M_{top} and ground.

This system is modeled as shown in the bond graph in Figure A-6. Input is modeled as a non-ideal flow source torquing the screw about the nut - essentially this is the artifact of nut motion that presents the risk of resonance addressed here. Transformer elements convert to linear force at the two endpoints; these transformers are dependent on the position of the nut z .

The sole input, τ_{in} , and the state variables are also shown in Figure A-6. This




Figure A-6: Bond graph representation of tuned mass-damper design for this system.

system produces the following state equations:

$$\begin{bmatrix} \dot{\omega}_i \\ v_{top} \\ F_{k1} \\ F_{k2} \\ F_{k3} \end{bmatrix} = \begin{bmatrix} \left(\frac{-b_1(l_1-z)}{I_{cm_screw}} - \frac{b_2(l_1+z)^2}{I_{cm_screw}} \right) & \frac{(l_1-z)}{I_{cm_screw}} & \frac{-(l_1-z)}{I_{cm_screw}} & \frac{-(l_1+z)}{I_{cm_screw}} & 0 \\ \frac{b_1(l_1-z)}{M_{top}} & \frac{-b_1-b_3}{M_{top}} & \frac{1}{M_{top}} & 0 & \frac{-1}{M_{top}} \\ k_1(l_1-z) & -k_1 & 0 & 0 & 0 \\ k_2(l_1+z) & 0 & 0 & 0 & 0 \\ 0 & k_3 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \omega_i \\ v_{top} \\ F_{k1} \\ F_{k2} \\ F_{k3} \end{bmatrix} + \begin{bmatrix} \frac{1}{I_{cm_screw}} \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \tau_{in} \tag{A.3}$$

This system was simulated in Matlab and the frequency response was examined. The frequency response for the function $\frac{\omega_i}{\tau_{in}}$ is shown in Figure A-7. The parameters used in this simulation included the mass values used in the preceding sections, with $k_1 = k_2 = k_3 = 3 \times 10^6$ N/m and $b_1 = b_2 = b_3 = 0$. This simulation also assumes that




Figure A-7: Magnitude and phase versus frequency for system with tuned vibration absorber. Input velocity over input torque.

the nut is at the center of the screw, so that $z = 0$. The desired mode appears as a dip in the magnitude response at a frequency close to 200 rad/sec. This is surrounded by two peaks in the magnitude response corresponding to the resonant frequencies of each of the isolated spring-mass modes. All three of these frequencies can be shifted around by changing the stiffnesses k_1 , k_2 and k_3 .

The problem with this approach for this application is that only the resonance at one specific frequency is cancelled. The effect is only useful here if it appears at all frequencies in the operating range. Since all system natural frequencies vary with position of the nut, while the driving frequency varies with the velocity of the nut, it is impossible to keep the resonance-cancelling mode aligned with the driving frequency, without introducing significant mechanical complexity. This task might, in fact, require an actively controlled mechanism. This degree of complexity is not desired at this time.

A.5 Conclusions

This analysis shows that the implementation of a tuned vibration absorber to eliminate any resonance problems introduced by a compliant mounting is prohibitively complex. Instead it is best to use the structure of Case 2 above, introducing compliance in the top or bottom bearing mount, or both, and maintaining a minimum stiffness such that the minimum resonant frequency is outside the system operating range.

Appendix B

Software Timing Study

It is useful to understand the approximate timing of operations performed in the software each sample period. This information can help in studying delay, in estimating a maximum possible sampling rate, and in optimizing the code. The results presented are for a QNX system running C++ on a 450 MHz Pentium II PC with 128 MB of RAM, and must be adjusted for different machines.

B.1 Testing

Some simple tests were run on the software to check the timing of various operations. To test the timing of a specific function or set of functions, a QNX timer loop was first set up. This timer was made to trigger an interrupt after some period of time, for example 1 sec. A loop was set up to begin after the timer started, and continue looping until the interrupt signaling the end of the timer interval was received. A counter was incremented each time through the loop. The approximate time for each trip through the loop was calculated by dividing the timer period by the number of times through the loop. The overhead in running the loop itself, computed by running the test with an empty loop, was subtracted from the total time to give the approximate time that the commands in the loop took to execute.

Several tests that involved negotiating through a library structure suggested that approximately 3 μ sec might be used each time the code must jump to another object.

This value was used in estimating the overall timing of the code.

It should be noted that the time to complete individual commands does not necessarily add linearly when several commands are run in sequence. For example, function A was timed at 18 μsec using the above experimental procedure. Function B was timed at 13 μsec in the same way. The experiment was then repeated with both function A and function B in the loop each time through. The time to run both functions was measured at 22 μsec . Because of this inconsistency, the results presented here must be understood to be only estimates.

Table B.1 gives a full list of the software processes that run between the beginning of a sample period and the output to the actuator, along with estimates of the time for each operation. The testing was done on code used for a two-link manipulator, so it involves some calculations that are not required for the screw module. Again, the times given are machine-specific.

B.2 Intentional Delays

In working with device drivers and other code that operates at an extremely low level of software, it is often useful to intentionally introduce delays. There are several different methods of doing so in QNX, and this section addresses these different approaches. Again, all testing was done on a Pentium II.

One way to introduce delay is to add a loop with a large number of bit-level operations. For example, a loop that incremented, then decremented an integer 100 times each time it ran added a delay of about 0.5 μsec . This method also seemed fairly linear; 200 times for each loop resulted in a 1.0 μsec delay, and so on. It is not a good idea, however, to use bit operations for this purpose. Since these operations are at the lowest level, the time to execute them is particularly sensitive to the operating conditions of the PC. While the delays were fairly consistent within several minutes, after several hours the test was repeated and produced results that differed by as much as 50% from the original tests.

A second method of adding delay is to use the `outp` command to write to a

dummy ISA address. This delay tested at about 1.4 μ sec, and increased linearly with the number of operations.

A much larger delay is introduced by sending information to the visual display, with either `cout` or `printf`. Each of these operations took about 330 μ sec to complete. This delay was also linear with the number of times the function was called.

For delays much greater than those discussed above, one might use the `delay()` command. This command takes an integer number of milliseconds as an argument. The delay is accurate above about 100 msec; below this value the results are highly irregular. A commanded delay of 1 msec seems to produce about a 10 msec delay.

Clearly, interfacing to hardware such as the graphics card or the ISA bus is the most effective way of introducing substantial delays. C and C++ are very fast with floating point and bit operations, so attempting to slow the code down with these tactics is largely ineffective.

Process	Time (μ s)	Elapsed (μ s)
clock.Wait() activities	2.21	
jump to Manta(): Update Sensors	3.0	
check clock_mode	0.01	
jump to uei_adc::Scan	3.0	
adc1.Scan, except for Done	7.46	15.68
jump to adc::Done	3.0	
jump to sensor::notify	3.0	
mapping in sensor notify	0.25	
jump back to update sensors	3.0	
jump to F.scan	3.0	
ATI_Read_Forces	128.0	155.93
jump back to update sensors	3.0	
float operations	0.1	
point2d formed	0.1	
create Jacobian	0.43	
point2d created, call J and point2d	0.2	
jump to F.resolve	3.0	
F.resolve()	0.05	
jump back to update sensors	3.0	
point2d with several J fcns run	1.0	
jump into J fcns	3.0	
jump out of J fcns	3.0	
point2d xy - mapping point2d angles	1.0	
jump to point2d mem fcns	3.0	
jump back to update sensors	3.0	
jump to point2d mem fcns	3.0	
jump back to update sensors	3.0	
jump to point2d mem fcns	3.0	
jump back to update sensors	3.0	
jump to point2d mem fcns	3.0	
jump back to update sensors	3.0	
jump to point2d mem fcns	3.0	
jump back to update sensors	3.0	
jump to point2d mem fcns	3.0	
jump back to update sensors	3.0	
total time to compute above fcns	1.0	
jump to system.state() - ctrlr call	3.0	
jump to clock.time() - ctrlr call	3.0	
jump to zctrlr	3.0	219.81
vector float operation	0.02	
loop to assemble x, v array	0.02	
vector float operation	0.02	
jump to manta()::jj	3.0	
point2d with Jt, point2d vector	0.2	
jump to fabs	3.0	
jump to fabs	3.0	
check for max torque	0.01	
vector operations	0.02	
jump to actuator::jj	3.0	
calculations	0.01	
jump to dac.output	3.0	
run dac.output	6.5	
jump back to actuator::jj	3.0	
jump to dac.output	3.0	
run dac.output	6.5	254.11

Table B.1: Approximate timing of software operations in each sampling period.

Appendix C

Models for Frequency Response

The frequency response data in Chapter 4 was presented along with the output of several fitted models that illustrated the sources of the main features of the experimental data. This appendix provides more detail about the models used in these simulations.

C.1 Simulated system

The system simulated is shown in Figure C-1. This block diagram shows the different parts of the models used for both the amplifier and actuator characterizations. The model for amplifier characterization that produced the results shown in Figure 4-8 includes the path from I_{in} to I_{out} shown in Figure C-1. The model for actuator characterization that produced Figure 4-10 is represented by the path from I_{in} to τ_{out} .

C.2 Amplifier Characterization Model

For the $\frac{I_{out}}{I_{in}}$ relationship, the current is subject to a delay equivalent to the sum of the amplifier delay and the software delay. The signal is then passed through a zero-order hold, equivalent to the D/A action of the PC controller. The signal passes through the amplifier, which is modeled as a first order butterworth filter with a cutoff frequency

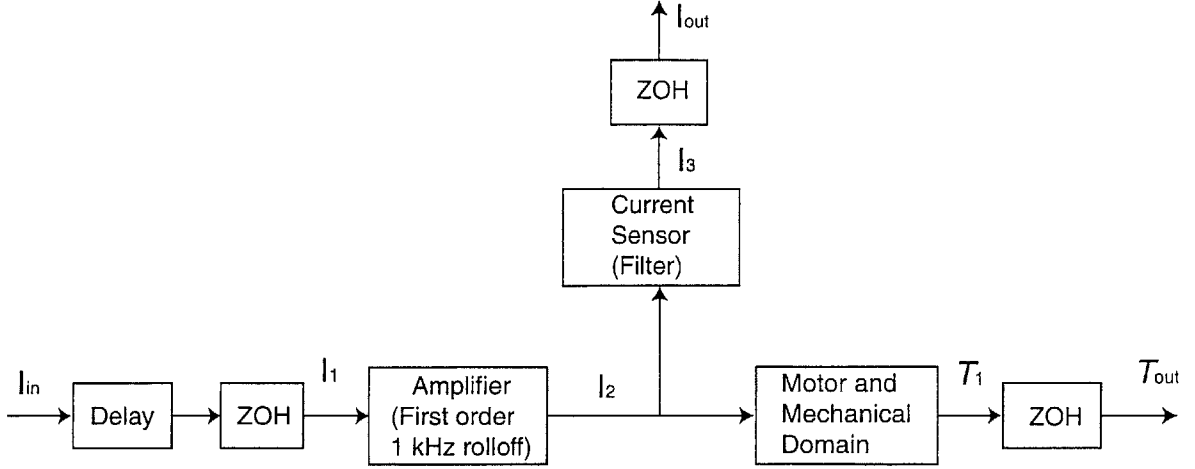


Figure C-1: Block diagram of models used in simulations to fit frequency response data in Chapter 4.

of 1 kHz. More specifically, this block of the diagram is represented by the following transfer function:

$$\frac{I_2}{I_1} = \frac{6283.2}{s + 6283.2} \quad (\text{C.1})$$

This output goes to the current sensor, which is primarily a low pass filter with a gain. The frequency response is determined by the circuit parameters. This block is represented by the transfer function:

$$\frac{I_3}{I_2} = \frac{R_s(1 + \frac{R_4}{R_3})\omega^2}{s^2 + (R_1 + R_2)C_1\omega^2s + \omega^2} \quad (\text{C.2})$$

where

$$\begin{aligned} R_s &= 0.1\Omega \\ R_1 &= 19600\Omega \\ R_2 &= 20000\Omega \\ R_3 &= 1000\Omega \\ R_4 &= 33200\Omega \\ C_1 &= 5 \times 10^{-9}F \end{aligned}$$

$$C_2 = 1.5 \times 10^{-8} F$$

$$\omega^2 = \frac{1}{R_1 R_2 C_1 C_2}$$

I_3 passes through another zero-order hold, representing the PC A/D data sampling. This produces I_{out} , the model equivalent of the data measured in the experiment.

C.3 Actuator Characterization Model

For the actuator testing, the same transfer function is used from I_{in} to I_2 . The current sensor, however, is removed from the model, as it does not play a role here. Instead the current runs into the block titled “Motor and Mechanical Domain”. This block represents the transition from current to torque as well as the dynamics in the mechanical system. The transition from current to torque is simply modeled as multiplication by a constant gain K_t . This constant is simply $K_t = 1$ here, because the software has already provided a mapping from desired torque to actual current. The transfer function for the block is:

$$\frac{\tau_{out}}{I_2} = \frac{K_t K}{s^2 + \frac{B}{J_m} s + \frac{K}{J_m}} \quad (C.3)$$

where K is the stiffness of the force transducer, J_m is the motor rotor inertia, and B is the damping value adjusted to fit the model. For the simulation that produced Figure 4-10, the values used for these parameters were:

$$K = 10000 \frac{Nm}{rad}$$

$$J_m = 1.96 \times 10^{-5} kg \times m^2$$

$$B = 7.2 \frac{Nm}{rad/sec}$$

The computed torque τ_1 is passed through a zero-order hold, representing the A/D card. The amplifier model is based on simulations of the model represented by $\frac{\tau_{out}}{I_{in}}$.

The models were simulated in Matlab and the results are presented in Chapter 4.

Appendix D

Conference Paper Detailing the Design

DETC2000/MECH-14151

3-D EXTENSION FOR MIT-MANUS: A ROBOT-AIDED NEURO-REHABILITATION WORKSTATION

Hermano I. Krebs¹, Stephen P. Buerger¹, Kristin A. Jugenheimer¹, Dustin Williams¹, Neville Hogan^{1,2}

¹ Massachusetts Institute of Technology, Mechanical Engineering Department
Newman Laboratory for Biomechanics and Human Rehabilitation

² Massachusetts Institute of Technology, Brain and Cognitive Sciences Department
77 Massachusetts Ave., 3-137

Cambridge, Massachusetts 02139

Phone: (617)253-8112; Fax: (617)258-7018; Email: hikrebs@mit.edu

ABSTRACT

An extension to an existing planar manipulator for robot-aided therapy is described. The MIT-MANUS rehabilitation workstation, shown to be effective in sensorimotor training of recovering stroke patients, is augmented by the addition of a module permitting movements in a 3-dimensional workspace. Key analyses in the mechanical design of the new module are described. A compact, screw-driven, low mechanical impedance module meets the target requirements. Manufacturing of this novel device will be followed shortly by its deployment in clinical trials.

INTRODUCTION

Rather than using robotics as an assistive technology for a disabled individual (a sophisticated version of a crutch), our research focus is on the development and application of robotics as a therapy aid. We envision robots and computers as supporting and enhancing the productivity of clinicians in their efforts to facilitate a disabled individual's functional recovery. To that end, we deployed our first robot, MIT-MANUS (see figure 1), at the Burke Rehabilitation Hospital, White Plains, NY in 1994. MIT-MANUS has been in daily operation for over 5 years, delivering therapy to over 100 stroke patients.




Fig.1. A recovering stroke patient receiving upper extremity robotic therapy with MIT-MANUS.

Most of the work accomplished so far has been aimed at the prevailing view that nature, e.g., lesion size and foci, are the overwhelming factors determining the stroke patient outcome and that nurture bestows no significant effect over the final outcome. If that were true, the application of robotics as "rehabilitators"¹ would be fundamentally flawed. Therefore, the fundamental question addressed in the last five years was whether exercising therapy influences brain recovery.

¹ To our knowledge, Steven Lehman (University of California, Berkeley) coined the term "rehabilitators."

CLINICAL RESULTS

As background, we will revisit some of the results of the initial clinical trial of MIT-MANUS (for more details see Aisen et al., 1997, Krebs et al., 1998, Volpe et al., 1999). Yet, to understand the implications of our results, suffice it to say that we have strong evidence that our robot-aids are not limited to entertainment. Our results suggest that nurture indeed has an effect on nature, and to exercise a hemiparetic or hemiplegic limb appears to harness and promote the process of neuromotor recovery following a stroke.

Twenty stroke patients exhibiting a unilateral lesion were enrolled in the initial clinical trial. Patients were typically enrolled 4 weeks following the onset of their stroke. Patients were randomly assigned to an experimental and a control group. The experimental group received an hour per day of robot-aided therapy, which consisted of a series of video games exercising the shoulder and elbow. If the patient could not perform the exercise routine, the robot guided or assisted the patient movements. The control group received an hour per week of "sham" robot-aided therapy with the same video games. If the patient could not perform the exercise routine, he/she used the unimpaired arm to guide to impaired one to complete the game's goal. The clinical team and the patients were kept blinded to the group assignment. Both groups were evaluated by the same blinded clinician (double blinded study).

Table I summarizes the results of the initial study as measured by standard clinical instruments: the upper limb subsection of the Fugl-Meyer (F-M), Motor Power for shoulder and elbow (MP), Motor Status Score for shoulder and elbow (MS1), and Motor Status Score for wrist and fingers (MS2).

Group	F-M (/66)	MP (/20)	MS1 (/40)	MS2 (/42)
	Δ 1	Δ 1	Δ 1*	Δ 1
RT	14.1	3.9	9.4	5.5
ST	9.9	2.3	0.8	4

Table I. Change during Acute Rehabilitation (20 patients): Experimental (RT) vs. Control (ST) Group - Δ 1 admission to discharge of rehabilitation hospital; one-way t-test that RT > ST with $p < 0.05$ for statistical significance (*).

Note that the MS1 for shoulder and elbow (focus of the exercise routines) of the experimental group shows a statistically significant difference over the control group in a comparable period. Note also that we observed no difference between groups in the MS2 for wrist and fingers (not exercised). This result suggests a local effect with limited generalization of the benefits to the unexercised limb or muscle groups. If this is the case, we must extend our robot-aids to exercise different groups of muscles and limb segments. Therefore we are presently designing and building

rehabilitators to work with different limb segments, e.g., wrist, fingers, legs, and to work with different group muscles, e.g., spatial arm movements. In this paper, we will describe the design of an additional module to the existing planar MIT-MANUS.

MIT-MANUS

To determine if exercise therapy influences the process of neuro-recovery following a stroke, we required a tool for controlling the amount of therapy delivered to a patient, while objectively measuring patient's performance. While robotics can provide such a tool, a commercial industrial robot will not suffice for this application. The typical industrial robot is intrinsically a position-controlled machine that does not readily yield to external forces. Active force feedback is needed to make the robot respond to the patients' actions, but that approach cannot (Lawrence, 1988) produce the "back-drivability" (low mechanical impedance) required to move smoothly and rapidly to comply with the patients' actions.

To address the limitation of commercial robots, we began in 1989 to design and build a novel low-impedance robot specifically for clinical neurological applications capable of interacting safely and gently with humans: MIT-MANUS (Hogan et al., 1995). In contrast to industrial robots, MIT-MANUS is configured for safe, stable and compliant operation in close physical contact with humans. This is achieved using impedance control, a key feature of the robot control system. Its computer control system modulates the robot's reaction to mechanical perturbation from a patient or clinician, ensuring gentle, compliant behavior. Impedance control (Hogan, 1985) has been the central contribution of Hogan's engineering research since the early eighties and has been extensively adopted by other robotics researchers concerned with human-machine interaction (see, e.g., the February 1997 issue of IEEE Control System, Special Issue on Robotics, which contains several articles on impedance control). MIT-MANUS is a planar robot that can move, guide or perturb the movement of a patient's upper limb and can record motions and mechanical quantities such as the position, velocity, and forces applied. MIT-MANUS design specification envelope was as follows:

General specifications for the target robot:

- The robot must be back-driveable.
- The robot should be modular, allowing expansion and combinations of actuator packages.
- Actuators are to be of electromagnetic type.
- Controller platform will be digital - PC based.
- All the sensors will be physically located in the drive package, i.e. position, velocity, and torque sensor.
- Working surfaces should be kept free and uncluttered to allow maximum visibility and allow easy access to the users, i.e. patients and clinicians.

- The robot should be portable and its weight targeted to 311N.
- The robot degrees of freedom will be limited in the present unit to the plane.

End-point specifications for the target robot:

- Workspace area of 0.381x0.457m.
- End-point static force of 45N.
- Motion control bandwidth of 1 Hz for a 0.152m motion
- End-point stiffness of 0 to 2 N/mm.
- From DC to 7 Hz the end-point impedance should behave primarily as viscous friction
- End-point inertia and anisotropy should be within 2/3Kg < I < 4/3Kg.
- Static friction and anisotropy should be within 56.7grams < Fs < 113.4 grams.

MIT-MANUS was completed in 1991. An overview of the robot main characteristics can be found elsewhere (Krebs et al., 1998).

SPECIFICATION FOR AN ADDITIONAL DEGREE OF FREEDOM

In the previous section we described the basic design specifications of the existing planar robot: MIT-MANUS. While it was tempting to limit the new design to be an extension of the existing one, for example, by attaching it as a module to the existing robot end-effector, we did not impose any such constraint. The new design could represent a substantial departure from the existing planar robot. Therefore we consider the same design specifications for the horizontal plane and in addition the following requirements for the vertical degree-of-freedom:

- Vertical Range of Motion – our clinical collaborators at the Burke Rehabilitation Hospital suggested that to be clinically relevant the new design should accommodate an 0.457m vertical range of motion. This is particularly important in ADL tasks (activities of daily living) such as feeding and grooming.
- Low Endpoint Friction – as explained in the previous section, our goal is to design a back-driveable (low mechanical impedance) robot. Therefore we set to minimize friction to the range of 113.4 to 226.8 grams.
- Low Endpoint Inertia – the end-point inertia and anisotropy should also be minimized. The maximum inertia should not exceed three times the maximum design specification of the planar robot, i.e., 4kg (horizontal plane) and 3/4kg for vertical movements.
- Ability to Withstand Non-Axial Loads – the human might be able to push up to 45N, therefore the device should be able to withstand up to 112.5N.

- To guarantee smooth operation, the drive system should minimize any source of backlash, which should not exceed 0.2mm.
- Ability to Vary Wrist Configuration -- the device's end-effector should permit adjustments to different wrist interfaces and configurations (e.g., wrist pronated or supinated).

POTENTIAL CONFIGURATIONS

To design a robot capable of delivering robot-aided therapy within a workspace volume of 0.381x0.457x0.457m, we considered two basic kinds of formulations: a) an integral design in which all three degrees of freedom are part of the same package, and b) a modular design in which the vertical motion is provided by a module mounted at the tip of existing planar mechanism. Both categories require that the vertical motion blend seamlessly with the planar motion, assuring the same "natural" experience as the patient moves his or her hand vertically, horizontally, or in space. In order to meet these specifications, a number of potential configurations were systematically generated considering all available modes for actuation and transmission. The resulting concepts can be loosely classified as follows:

Integral configuration

Three basic configurations of this kind were considered: a) a SCARA configuration spanning the vertical plane mounted on a rotational base; b) a configuration similar to the existing planar robot in which the whole set is moved vertically at the base; and c) a comparable configuration for which the actuators' location remains stationary, but a linear motor slides the SCARA arms up and down two keyed shafts. The primary disadvantage of this class of designs is in the sizing of the actuators, which have to carry either the weight of others or of the arm.

Vertical module at the end of the SCARA with transmission

To minimize the actuators' size, we considered the alternative concept of adding a vertical module to the tip of the existing planar robot driving from the base through a transmission line. Transmissions included timing belts and bevel gears. The mechanism per se transformed rotary into linear movement via a rack-and-pinion, lead screw, or friction drive. The primary disadvantage of this class of design is in the dynamics introduced by the transmission line.

Vertical module including actuator at the end of the SCARA

To minimize any transmission line effects, we considered configurations with the actuator located in the vertical module at the tip of existing planar robot. Different designs were examined including a rotary motor driving a lead screw or a rack-and-pinion, a slider-crank mechanism, and a direct drive linear motor. Despite the gains in efficiency with such

design, this class is handicapped by the weight of the actuator, which increases planar inertia, as the actuator must be moved with the endpoint.

In the next section, we will discuss the merits of each of the configuration, and cite the most promising configurations.

SEMI-FINALIST CONFIGURATIONS

To judge the merits of each configuration, we graded them against a performance index composed of the specifications and their corresponding weighting factors. Friction and inertia were the heaviest coefficients and the favored configuration consisted of a modular design mounted at the tip of the existing planar robot driven by a brushless DC motor.

Of the other configurations, note that moving or rotating the entire robot would substantially increase inertia and require larger actuators. "Binding" when back-driving limited the sliding SCARA-arm design. Linear motors at the tip of the planar robot were excluded on the basis of inertia, since an appropriate choice would weigh at least 112N. Rack and pinion assemblies are bulkier and present more backlash than lead screws. Friction drives do not tolerate non-axial loads well. The lead screw appeared to offer advantages in the most critical areas, i.e., a properly selected screw has very high efficiency and low inertia, and is easily back-driven. While the screw was the best option for the vertical device, two configurations did score equally well. They differ primarily by the location of the DC motor and the existence of a transmission line in one of them.

The self-enclosed option consisted of a brushless DC motor mounted to the screw at the endpoint of MIT-MANUS. Such a design would be compact, simple, and self-contained. On the other hand, it could result in a significant increase in inertia, since the planar configuration must carry the vertical actuator. The alternative was to install the motor at the base of the robot and to use a timing-belt contained inside the SCARA arms to convey power to the lead screw at the tip of the SCARA. While reducing the system deadweight and increasing the potential payload, it introduced the dynamics of the timing-belt.

In what follows, we present some results on the search for the optimal alternative. In particular, we analyzed the effects and anisotropy introduced by additional inertia at the tip of the planar mechanism, the low resonant frequency of the belt driven transmission, and the selection of screw characteristics and materials.

BELT-DRIVE

To recap, in an attempt to minimize added endpoint inertia, we considered keeping the vertical actuator at the base driving the lead screw through a timing belt. The major concern in using a belt drive system is resonance in the system: resonance between the sides of the belt (the "guitar

string" effect) and between the two pulleys due to compliance in the belt and compliance between the belt and the pulley teeth. The latter type can occur at a low frequency, which would interfere with the system control. To check whether such vibration would be a problem in the system, consider the following simplified linear system of two rotary inertias (the motor and the load) with a rotational spring and dashpot between them (the belt). Both inertias were coupled to ground by rotational dashpots representing friction. The motor inertia was coupled to a torque source representing the motor.




Fig.2. Belt-Drive Model

The derived state equations used were (where the index m stands for motor and l for load with $\Delta\Theta$ representing the angle between J_m and J_l):

$$J_m \omega_m = T_s - B_m \omega_m - K_\phi \Delta\Theta - B_\phi (\omega_m - \omega_l) \quad (1)$$

$$J_l \omega_l = K_\phi \Delta\Theta + B_\phi (\omega_m - \omega_l) - B_l \omega_l \quad (2)$$

$$\frac{d\Delta\Theta}{dt} = \omega_m - \omega_l \quad (3)$$

We estimated the parameters for a sensible configuration using off-the-shelf components, i.e., inertia, ground friction, belt stiffness and damping coefficients, and simulated the model state-space representation in three separate cases: a) no damping, b) only belt damping, and c) entire system. The simulation shows that in the no damping case, the system was marginally stable with resonance at 3 Hz and a bandwidth of 10.6 Hz. When the belt damping was included in our simulation, no resonance was observed and the system bandwidth increased to 51.7 Hz. Lowering the damping coefficient of the belt by a factor of ten uncovered again the 3-Hz resonance with a damping factor of 0.446. Finally, simulating the whole system reduced the bandwidth to 40 Hz, but as in the previous case (only the belt damping) we observed no resonance.

Injecting belt damping extends the system bandwidth (belt damping case), which physically indicates that the system behaves more like a rigid body. On the other hand, injecting damping to ground reduces the system bandwidth (whole system case), which is expected for a mass-spring-damper case. In summary, the whole system bandwidth of 40 Hz exceeds our requirements (see system specifications), although the potential appearance of the natural frequency at

Hz is of concern. The belt-drive configuration remains as a potential implementation.

SELF-ENCLOSED DRIVE

The addition of the vertical module at the tip of the planar mechanism will alter the effective endpoint inertia. We estimated the weight of the new module composed by a DC motor directly driving a lead screw resting on a rail to be under 2.7Kg. Originally we specified a maximum anisotropy of 2:1 for the planar mechanism. Our estimates suggest that the final planar mechanism is comfortably below this limit with the inertia anisotropy ranging between 1.23 to 1.53 (ratios between major and minor axes of the inertia ellipsoid matrix at different workspace locations). Adding the estimated vertical module weight to tip of the planar mechanism further reduces this anisotropy. Therefore the critical constraint to the self-enclosed drive configuration is the maximum value of the major axis of the inertia ellipsoid, as well as the vertical inertia. Our estimates suggest that for the worst location it stays within the 4kg maximum inertia (see figure 3) and the vertical inertia is limited to 0.64kg. Both estimates reside within our initial requirements.




Figure 3. Estimate of End-Point Inertia Ellipsoids with the additional Vertical Module

Summarizing both modular configurations, i.e., self-enclosed and belt driven, are feasible configurations scoring well against the requirements. We decided in favor of the vertical self-enclosed drive module primarily due to its simplicity and self-contained structure. The relative complexity of the belt-drive implementation turned it into our second best alternative.

SCREW DRIVE

As previously stated, the lead screw appeared to offer advantages in the most critical areas, i.e., a properly selected

screw has very high efficiency and low inertia, and is easily back-driven. In what follows, we will discuss the selection of an appropriate screw and lead (the linear nut advancement per screw revolution). A smaller lead permits the drive motor to run at a faster speed, and therefore a lower torque, resulting in the use of a smaller motor and reduced end-point inertia. Small leads are the most common in driven power screw applications. On the other hand, to guarantee back-drivability, it is advantageous to maintain a large lead (a box will slide down more quickly on a steeper ramp). There is one additional factor to be considered in the tradeoff between actuator inertia and system back-drivability: friction. Reducing screw friction also enhances back-drivability.

Using a classical screw model, the thread is unwrapped from the screw and the system can be modeled as a block (the nut) on an inclined plane (the screw thread). The system is shown in Figure 4.a. If the nut is removed and replaced with a normal force N and a friction force (whose direction is dependent on the forces F and P) the model looks like Figure 4.b. F is the input force driving the nut down; P is the force from the motor torque attempting to drive the nut up. ϕ is the lead angle, and μ is the coefficient of friction. θ is a generalized thread angle, incorporating both thread angle and lead angle. At equilibrium we have:

$$P - \frac{F(\cos\theta * \tan\phi - \mu)}{\cos\theta + \mu * \tan\phi} = 0 \quad (4)$$

The proper statement of this equation depends on the causality of the forces. If F is the force in the direction of motion, it is opposing friction and therefore must be attached to the friction terms. For the F term to be greater than zero, a certain relationship must exist between ϕ , the lead angle, and μ , the coefficient of friction. Likewise for the case of the screw driving the nut up, we get another such relationship. Both are described below:

$$\tan\phi * \cos\theta > \mu \quad (5)$$

$$\frac{\cos\theta}{\tan\phi} > \mu \quad (6)$$

Both conditions must be satisfied to drive the system from either side. To estimate the influence of an order of magnitude change in the coulomb friction on the back-drivability, we used the same set of geometric values and masses with a plastic lead screw ($\mu=0.15$) and a roller screw ($\mu=0.015$). An order of magnitude reduction in the coulomb friction resulted in a 40% greater acceleration to a given input force for the roller screw. Roughly 30% of the input force was lost to friction with the plastic screw, while less than 5% was lost with the roller screw.




Fig.4. Nut sliding on a Wedge

This analysis suggests that minimizing friction might allow us to choose a smaller lead, which means smaller actuator inertia, for the same degree of back-drivability. Therefore we opted for a ball screw or a roller screw rather than a plastic lead screw; since ball screws tend to have inconsistent friction due to the balls rubbing against each other, we searched for commercial roller screws. We settled on a Flennor Rollnut, which has two conical rollers that ride in an angled track around the screw, theoretically in 100% rolling contact, with no sliding. Thus the only theoretical sources of screw friction are the bearings that hold the rollers and the bearings on the screw.

VERTICAL MODULE

The final configuration incorporates the rollnut and screw (Flennor Corp.) into a linear guide system with different handles for the patient. The rollnut is mounted to the carriage of a ball Monorail (Schneeberger Inc.). This carriage travels with four-point contact along a rail mounted to a U-channel support. The support is rigidly connected to the end of the last SCARA linkage. A brushless, slotless DC motor (Aerotech Inc.) is mounted atop the support to drive the screw. Bearings in the top and bottom support plates reduce friction and the load on the motor. The screw can be locked in position with a spring-loaded pin at the bottom of the structure. This allows the therapist to put the handle at a specific vertical position and then have the patient exercise only in that horizontal plane. Note the bracket mounted to the rollnut. Different interfaces can be attached to this bracket, enabling the attachment of a new interface with rotation in any direction. This new handle allows the therapist to move and lock it to any desired angle for the duration of the exercises. The final configuration is illustrated in Figure 5.

CONCLUSION

It is not far-fetched to suggest that robot-aided neuro-rehabilitation might confer a genuinely positive effect on the process of neuro-recovery following a stroke. The pioneering clinical results of our prior work are consistent with a prominent theme of current neuroscience research into the sequelae of brain injury or trauma, which posits that activity-dependent plasticity underlies neuro-recovery. Furthermore, our results with more than 100 stroke patients open up a

number of opportunities for the robotics community. We envision the rehabilitation clinic of the future as gyms of rehabilitators working with different limb segments, muscle groups, and functional tasks. At this gym, the therapist tailors an exercise routine to the particular patient needs, optimizing his/her recovery, while increasing productivity by overseeing several patients at the same time each working with a robot-aid. We also envision rehabilitators at patients' homes and the increase of productivity of the overall rehabilitation system via robotics quantitative measurement capabilities, and the potential reduction in "paper trail." MIT-MANUS was the first step in making this vision a reality.

From the realm of science fiction to the substance of humbling reality, the novel module for vertical movement is another marker along the trail. The module is currently in the final stages of manufacturing and we expect to deploy it at the Burke Rehabilitation Hospital prior to end of this summer. It follows the same guidelines of MIT-MANUS, which includes back-drivability. Our experience proved that this feature is essential for any successful robot-aid and it represents a departure from commercial industrial robots, which are intrinsically position-controlled machines. Finally, while very little technology presently exists or is being developed to support the recovery phase of rehabilitation, we believe the landscape will change quickly in the near future, and to that extent we are designing other robot-aids for additional limb segments and muscle groups.




Fig.5. Illustration of the Vertical Module

ACKNOWLEDGMENTS

This work was supported in part by The Burke Medical Research Institute and NIH Grant # 1 R01-HD37397-01.

REFERENCES

1. Aisen, M.L., Krebs, H.I., McDowell, F., Hogan, N., Volpe, N., The effect of robot assisted therapy and rehabilitative training on motor recovery following stroke. *Archives of Neurology*, 54(1997), pp.443-446.
2. Krebs, H.I., Hogan, N., Aisen, M.L., Volpe, B.T., Robot-aided neuro-rehabilitation, *IEEE Trans. on Rehab. Eng.*, 6:1(1998), pp.75-87.
3. Volpe, B.T., Krebs, H.I., Hogan, N., Edelstein, L., Diels C., Aisen, M.L., Robot Training Enhanced Motor Outcome in Patients With Stroke Maintained over Three Years, *Neurology*, 53(1999), pp.1874-1876.
4. Lawrence, D. A. Impedance control stability properties in common implementations. *Proc. IEEE Int. Conf. Robotics & Automation*, (1988), pp.1185-1191.
5. Hogan, N., Krebs, H.I., Sharon, A., Charnnarong, J., Interactive robotic therapist, *U.S. Patent # 5,466,213*, *Massachusetts Institute of Technology*, (1995).
6. Hogan, N. Impedance Control: An approach to Manipulation: Part I, Part II, Part III. *J. Dynam. Syst., Measurement, Control – Trans. ASME*, 107(1985), pp.1-24.

Bibliography

- [1] Adebiyi, D. (1998) Fabrication and Characterization of Beta-Prototype MIT Manus: an Intelligent Machine for Upper-Limb Physical Therapy. MSME Thesis, Massachusetts Institute of Technology.
- [2] Aisen, M.L., Krebs, H.I., Hogan, N., McDowell, F. and Volpe, B.T. (1997) The Effect of Robot-Assisted Therapy and Rehabilitative Training on Motor Recovery Following Stroke. *Archives of Neurology* **54**:443-446.
- [3] Armstrong-Helouvry, B. (1991) Control of Machines with Friction. Kluwer Academic Publishers, Boston.
- [4] Armstrong-Helouvry, B., DuPont, P., Canudas de Wit, C. (1994) A Survey of Models, Analysis Tools, and Compensation Methods for the Control of Machines with Friction. *Automatica* **50**:1083-1138.
- [5] Bauchau, O.A. (1999) On the Modeling of Friction and Rolling in Flexible Multi-Body Systems. *Multibody System Dynamics* **3**:209-239.
- [6] Canudas de Wit, C., Olsson, H., Astrom, K.J., Lischinsky, P. (1995) A New Model for Control of Systems with Friction. *IEEE Transactions on Automatic Control* **40**:419-425.
- [7] Charnnarong, J. (1991) The Design of an Intelligent Machine for Upper-Limb Physical Therapy. MSME Thesis, Massachusetts Institute of Technology.
- [8] Colgate, J.E., Hogan, N. (1988) Robust Control of Dynamically Interacting Systems. *International Journal of Control* **48**:65-88.

- [9] Dahl, P.R. (1968) A Solid Friction Model. *The Aerospace Corporation, El Segundo, Calif.* TOR-158(3107-18).
- [10] Dahl, P.R. (1976) Solid Friction Damping of Mechanical Vibrations. *AIAA Journal* **14**:1675-1682.
- [11] Diffrient, Tilley, Harman. *Humanscale 7/8/9*. MIT Press, Cambridge: 1981.
- [12] Doeringer, J.A. (1999) An Investigation into the Discrete Nature of Human Arm Movements. Ph.D. Thesis, Massachusetts Institute of Technology.
- [13] Exlar Corporation. Product Catalog, 1999.
- [14] Flash, T. and Hogan, N. (1985) The Coordination of Arm Movements: An Experimentally Confirmed Mathematical Model. *Journal of Neuroscience* **5**: 1688-1703.
- [15] Foster, C. (1999) A Performance Characterization of an Interactive Robot. MSME Thesis, Massachusetts Institute of Technology.
- [16] Franklin, G.F., Powell, J.D., Workman, M. (1998) Digital Control of Dynamic Systems. Addison-Wesley, Reading, Massachusetts.
- [17] Gelinck, E.R.M., Schipper, D.J. (2000) Calculation of Stribeck curves for line contacts. *Tribology International* **33**:175-181.
- [18] Hess, D.P., Soom, A. (1990) Friction at a Lubricated Line Contact Operating at Oscillating Sliding Velocities. *ASME Journal of Tribology* **112**:147-152.
- [19] Hogan, N. (1985) Impedance Control: An Approach to Manipulation: Part 1 - Theory, Part 2 - Implementation, Part 3 - Applications. *Journal of Dynamic Systems, Measurement, and Control* **107**:1-23.
- [20] Kang, Y.-S., Kim, K.-J. (1997) Friction Identification in a Sight Stabilisation System at Low Velocities. *Mechanical Systems and Signal Processing* **II**:491-505.
- [21] Kerk Motion Products , Inc. Product Catalog, 1997.

- [22] Krebs, H.I., Buerger, S.P., Jugenheimer, K.A., Williams, D. and Hogan, N. (2000) 3-D Extension for MIT-Manus: A Robot-Aided Neurorehabilitation Workstation. *Proc. ASME IDETC/CIE*. September 10-13, 2000, Baltimore, MD.
- [23] National Stroke Association. Brain Attack Statistics. *www.stroke.org*, 2001.
- [24] Rao, S. (1995) Mechanical Vibrations. Addison-Wesley Publishing Company, New York.
- [25] Swevers, J., Al-Bender, F., Ganseman, C.G., Prajogo, T. (2000) An Integrated Friction Model Structure with Improved Presliding Behavior for Accurate Friction Compensation. *IEEE Transactions on Automatic Control* **45**:675-686.