Stable, High-Force, Low-Impedance Robotic Actuators for Human-Interactive Machines

by

Stephen Paul Buerger

S.M.M.E., Massachusetts Institute of Technology (2001)
B.M.E., University of Dayton (1999)

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

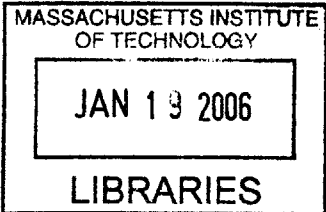Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2005

Author . . . . , ...
Department of Mechanical Engineering
February 11, 2005

Certified by . . . .
Neville Hogan
Professor
Thesis Supervisor

Accepted by . . . .
Lallit Anand
Chairman, Department Committee on Graduate Students

Stable, High-Force, Low-Impedance Robotic Actuators for Human-Interactive Machines

by

Stephen Paul Buerger

Submitted to the Department of Mechanical Engineering
on February 11, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Robots that engage in significant physical interaction with humans, such as robotic physical therapy aids, must exhibit desired mechanical endpoint impedance while simultaneously producing large forces. In most practical robot configurations, this requires actuators with high force-to-weight ratios and low intrinsic impedance. This thesis explores several approaches to improve the tradeoff between actuator force capacity, weight, and ability to produce desired impedance.

Existing actuators that render impedance accurately generally have poor force densities while those with high force densities often have high intrinsic impedance. Aggressive force feedback can reduce apparent endpoint impedance, but compromises coupled stability. The common standard for ensuring coupled stability, passivity, can limit performance severely. An alternative measure of coupled stability is proposed that uses limited knowledge of environment dynamics (e.g. a human limb) and applies robust stability tools to port functions. Because of structural differences between interaction control and servo control, classical single-input, single-output control tools cannot be directly applied for design. Instead, a search method is used to select controller parameters for an assumed structure. Simulations and experiments show that this new approach can be used to design a force-feedback controller for a robot actuator that improves performance, reduces conservatism, and maintains coupled stability.

Adding dynamics in series to change an actuator's physical behavior can also improve performance. The design tools developed for controller design are adapted to select parameters for physical series dynamics and the control system simultaneously. This design procedure is applied to both spring-damper and inertial series dynamics. Results show that both structures can be advantageous, and that the systematic design of hardware and control together can improve performance dramatically over prior work.

A remote transmission design is proposed to reduce actuator weight directly. This design uses a stationary direct-drive electromagnetic actuator and a passive, flexible hydraulic transmission with low intrinsic impedance, thereby utilizing the impedance-

rendering capabilities of direct-drive actuation and the force density of hydraulic actuation. The design, construction and characterization of a low-weight, low-friction prototype for a human arm therapy robot are discussed. Recommendations and tradeoffs are presented.

Thesis Supervisor: Neville Hogan
Title: Professor

Acknowledgments

Connections to friends outside of MIT have always been essential for me. Mike Burrows always reminds me of the importance of carefree days. Kirk Vashaw is a great friend and counterpart for heated political debates. No one is quicker to seize the opportunity to go out than Kevin O'Neill (who was never really an MIT guy). I also especially thank all the other Boston people, including Theresa & crew, Matt Cole, and all of Nicole's friends. People from elsewhere provided essential support and the chance to escape, whether on the phone or in person. Thanks to Patrick King, Patrick Schubel, Rob O'Leary, Matt Groves, Jill & Rob, Guapo, Casey, Chris Luckett, Purcell, Paul Haney, and many others, including my entire extended family on both sides.

I give a great big thanks to my parents for listening to my constant complaints and for continually assuring me that everything would work out in the end. I am sure they are glad this phase is over. Thanks to my sister Julia, her husband John and especially my two nieces, Emily and Rebecca, for reminding me that there are many more important things in life than grad school (for example, being sure that your noodles get exactly the right amount of sauce on top: none).

And finally I thank Nicole. No one has had a more unfiltered view of the challenges I've faced over the past few years and my reactions to them, and she has stood firm in the face of all my flaws longer than I ever could have hoped. Her caring smile has brought joy to the bleakest days. For that I owe her everything. Thank you Nicole.

Contents

List of Figures

13

14

15

17

18

19

List of Tables

Chapter 1

Introduction to the Interaction Problem

In popular culture, robots have a strong presence in the human realm. In television and movies, futuristic robots are servants and companions, living and working in close contact with humans. In reality, the actual use of robots has been mostly limited to factories and laboratories, tightly controlled environments where they can perform the repetitive, precise tasks at which they excel. The only humans who come into close contact with robots are those who design, build, test, and maintain them. Despite decades of development, robots have yet to assume a significant presence in the everyday lives of average citizens. Cost is undoubtedly one reason for this divide, but economies of scale could make many robots affordable. The more important reason is a technological barrier. Several large technical hurdles hinder both safety and performance and together prevent robots from being practical additions to daily human life. If a robot is to operate in a context in which it may physically contact humans, it must be assured that this contact cannot inflict injury. With today's technology, ensuring safety usually means sacrificing performance (e.g. by lowering force levels) to the point that the robot is ineffective at its primary tasks.

Within the last decade, specialized robotic devices have become slightly more common as applications have been identified in which robots can be successful without requiring potentially dangerous high forces. Robotic vacuum cleaners are a com-

mercial success [51]. Commercial haptic devices are becoming more affordable and practical for eventual home use [105]. These forays into the mass market are a result of a combination of demand, technical development, and clever choice of application. Both of the technologies cited above perform tasks (carpet sweeping and information exchange, respectively) that do not require high forces, and need not threaten human safety.

For robots to take a more significant role in human life, they must safely manage intentional and unintentional physical contact with humans, even while performing high-force tasks. Human interaction is not the only robot task that requires physical contact with an environment; many other machines face similar duties. For example, if a construction crane is required to move an object, it must physically contact that object (its environment) as it performs its duty. An assembly robot must be capable of moving parts and mating them with constraints (the challenging nature of these tasks is why robots are routinely used for non-contact manufacturing tasks like spot-welding and painting, but rarely for assembly). Robots and other machines (e.g. earthmovers, cranes, etc.) that must physically interact with their environments face a unique set of challenges in achieving both stability and performance. These challenges are brought to the fore when interacting with the environment produces coupled system dynamics that differ significantly from the dynamics of the robot system alone.

This thesis addresses the problems posed in achieving stability and performance in robots that interact with environments. In particular, the focus is on robots that are designed for direct physical contact with humans, but much of the work can be applied to other machines as well, when interaction with an environment, particularly one with well-understood mechanical properties, is required. Some of the proposed solutions may extend to all interactive robots; others are focused more directly on the human interaction problem.

Human-interactive robots

Several active areas of research concern robots that physically contact humans. Robotic aids for physical therapy [1, 62, 75, 18, 28] promise to alleviate therapists of the physical rigors of their work, freeing them to concentrate on patient needs. Robots excel at repetition and do not tire, so they can allow a patient to complete a greater number of moves in a therapy session, with improved consistency. Robots for physical therapy are discussed further below, and throughout this thesis.

Haptic devices [80, 10] provide a physical means of information exchange between humans and computers. Haptic devices simulate "virtual objects" at the human interface, so that the operator can gather information by touch. Unlike a conventional mouse, keyboard, or monitor, a haptic device is both an input and output device, permitting information to flow from user to computer and from computer to user. Haptic devices usually have low force capacity (i.e. 2-10 N), but increased forces could permit use in a broader set of applications.

Teleoperated master-slave systems [111, 66] are intended to permit human intervention in tasks that are too hot, too radioactive, too strenuous, or otherwise too dangerous for humans. They have gained prominence in the handling of nuclear materials, and can also be used to amplify human force or range of motion. A teleoperation system consists of a *master*, a device that the human operator contacts physically, a *slave* that emulates the motion of the master, and a controller. Usually the operator provides motion input and the slave moves accordingly, measuring the contact forces with its environment and feeding them back for scaled reflection by the master. Both the master and slave must have sensors and actuators, and both must contact environments. The requirements of the master are very similar to those for haptic devices, and the slave usually requires relatively high force capacity and the ability to stably interact with its environment, often consisting of objects to manipulate.

Human extenders [69, 70] are also intended to boost the force capacity of human limbs while retaining the human brain as the high-level controller. These devices exchange power and information with a human operator at one location and the

load at another. Unlike in master-slave systems, a portion of the load is directly transmitted to the human through the structure of the device, while the extender bears the rest of the load [70]. This eliminates the need for a second set of actuators to provide force reflection to the human operator.

Cobots [94, 46] represent an alternative approach for human-robot cooperation. Cobots help humans accurately manipulate heavy objects by rendering virtual surfaces. In contrast to haptic devices, however, the large forces needed to render virtual constraints come from passive physical phenomena rather than actuators. For instance, a cobot motor can position a wheel to create a virtual surface perpendicular to the wheel's axle. Motion parallel to this surface meets only the rolling resistance of the wheel; motion perpendicular to the surface requires overcoming *sliding* friction of the wheel. The actuator need only be strong enough to position the wheel, but the forces the surface provides may be much higher. One use for such a virtual surface is in helping a human guide a heavy object into place, for example during assembly, without colliding with nearby objects. More complex cobots can be made with continuously variable transmissions.

Robots are useful in several non-invasive, minimally-invasive, and fully invasive surgical applications [97]. Teleoperated systems are used for minimally invasive surgery [101]. In this case, the robotic system must interface with the hands of the human operator on one side and with human tissue on the other side. Force reflection can help a doctor recover some of the information that is lost with the ability to touch internal body structures directly, as is done in traditional surgery. In related work, robotic and haptic technology is used to simulate minimally invasive surgery to provide training [9].

Vehicle simulators [103] typically carry one or more humans as passengers to provide training in the operation of airplanes or heavy machinery. These systems move in response to the user's actions, as dictated by the control system. They interact with the human via information flow, by responding to the operator input, as well as via power flow, as they move the subject.

Powered prosthetic devices [12, 71] seek to restore limb function including active

intentional joint movement, fed by mechanical, electric, or neural inputs. This field is experiencing resurgent interest from the military. Powered prostheses are usually limited by force- and power-densities, making them too heavy to be useful. A substantial improvement in actuator technology could dramatically improve this field.

Exoskeletal robots [11, 82, 90, 117] have been proposed for some of the applications suggested above, such as physical therapy, and for other applications such as human performance enhancement. The weight of these devices, including the actuators, must either be borne by the human or connected to the ground. Because of this, lightweight actuators are desired, but desired force levels are comparable to human limb forces. Like prosthetics, these devices have been limited by actuator force densities.

Overall goal

The overall goal of this thesis is to improve the tradeoff between high force capacity, low mechanical endpoint impedance (or backdrivability), and low mass in actuators for devices that interact with humans. Many of the applications described in the previous section could benefit from an improvement in actuator technology, as they are limited by actuation capabilities. Changes are proposed to the design of both controllers and physical actuation systems to improve the aforementioned tradeoff, for example by reducing endpoint mass without reducing force capacity or increasing impedance. Modular and scalable solutions are sought so the technology can be used for a wide variety of applications. Physical interaction with the environment is the critical objective, and is specifically analyzed throughout.

1.1   Defining Mechanical Interaction

Throughout this thesis, the terms *mechanical interaction* or simply *interaction* are used to describe the task that distinguishes the machines addressed here from other types of robots and machinery. The definition of these terms is as follows:

> **Definition:** A robot *mechanically interacts* with its environment when the dynamics of the coupled system, consisting of the robot and environ-

ment in contact, differ significantly from the dynamics of the robot system alone.

The difference in dynamics can occur either in structure or parameter values, or both. Interaction can be detrimental both to stability and to traditional measures of robot performance, as is shown in detail in the next chapter.

The definition of interaction is intentionally qualitative. The quantitative significance of the interaction depends not only on the two systems and the way they are coupled, but also on the performance and stability requirements. If interaction affects the system in such a way that neither stability nor performance is compromised, it can generally be ignored. In many cases, however, particularly in almost all human-robot applications, interaction does affect performance, stability, or both, and must be analyzed.

1.1.1 Signal and limited physical interaction

Not all robots mechanically interact with their environment. In fact, the vast majority of robots used for industrial applications do not interact significantly. This is largely a consequence of the challenges that such robots pose; robots have gained prevalence mostly for applications that do not require interaction to be addressed specifically.

In many cases, a limited amount of physical interaction can be tolerated without specific modeling or control. If coupling to an environment does not produce appreciable change in the dynamics of a system, *signal interaction* may result, wherein the system passes information to the environment without effect on the source system. This concept is familiar to the analysis of electrical systems; when two systems with sufficiently mismatched impedance are connected, they can be analyzed separately, and the interaction between the two can be treated as signal flow. When the impedances of the two systems are similar, however, this approach fails. The analog between electrical impedance and mechanical systems is quite fruitful, and is further explored in chapter 2.

Similarly, limited physical interaction can sometimes be accommodated by mar-

gins of stability and performance. Interaction forces, for example, can be represented as disturbance inputs when a control system is designed, to ensure that stability and performance are achieved whether or not there is interaction. If interaction produces only relatively small changes in parameter values, then inclusion of uncertainty in controller design might permit the maintaining of stability and the achievement of performance objectives. In either case, if the interaction is significant then the disturbance or uncertainty grows too large, and the performance suffers mightily under the requirements for stability. This case requires directly addressing interaction by the methods described in the next chapter.

Signal interaction is analyzed thoroughly in the control literature; an assumption of signal interaction underlies most block-diagram analysis. Classical and modern control theory can also encompass limited mechanical interaction with disturbance or noise rejection, or with robust control [45, 79]. This thesis is concerned with cases where these techniques are not sufficient, and interaction is significant. Actuation and control techniques are studied to directly manage physical interaction. Many of these solutions apply to the broad interaction problem, but are here focused on the details of robots interacting physically with humans. The specific implications of this focus are discussed below.

1.2 Devices that Interact with Humans

This work focuses in particular on machines that mechanically interact with humans. Some examples are provided above. Each of these types of devices can require mechanical interaction with human limbs or tissue. For the purpose of clarifying the unique requirements that such devices face, it is useful to classify these devices on the basis of their functional differences.

1.2.1 Restrictive devices

Some human-interactive machines must impede or guide movement of the human, but need not themselves provide active motion. Such devices may be required to

33

create physical constraints or drain energy, but not to generate it. They can function as haptic devices or training aids for various tasks. Machines of this type can, in general, rely on brakes in lieu of actuators. An example of a configuration that uses magnetic particle brakes in a haptic device is described in [102]. Cobots [94, 46] are another implementation of this class of device. Cobots rely on manipulating physical constraints to allow motion in some directions but not in others, and do not require their actuators to directly produce targeted force levels.

1.2.2 Low-force active devices

Unlike the first class, devices in this class may be required to generate energy and initiate movement, as well as to dissipate and store energy, but are limited to relatively low saturation forces. These machines are used to produce tactile sensation, to exchange information, and to render virtual objects. Like restrictive devices, these devices must permit free motion under certain conditions. Almost all commercial haptic devices fit into this class, as they are limited to around 10 N or less. Examples of haptic devices include the PHANTOM [80] and the magnetically levitated wrist device [10].

1.2.3 High-force active devices

The most challenging interactive devices are required to generate, dissipate, and store energy, and to do all of this at high force levels. Some tasks may be quite similar to those of the previous class, including rendering virtual objects, but others may include moving a human limb or performing a task that the limb might otherwise perform. These devices must permit free motion at times, but maintain the capability to actively move against substantial forces. Actuators that are capable of producing high forces are, as a rule, more massive than actuators of similar design capable of producing smaller forces, and are also generally more massive than braking technologies that provide the same force levels; therefore achieving low impedance is more difficult for this class than for the previous classes. Human limbs are biological members of this

class; they are capable of high forces with low impedance, thanks to the actuation capabilities of the musculoskeletal system. In general, robots that interact with human limbs may have force and impedance specifications comparable to the capabilities of the limb with which they interact. Because technology is unavailable to match the function of human limbs, this presents an engineering challenge. Potential applications for devices that satisfy these requirements are plentiful. Attempts have been made in rehabilitation robotics [28, 62, 117], extenders [69, 70], assistive devices [82], powered prosthetics [71], and other areas. A central challenge of all of these designs is the tradeoff between actuator force density and impedance (discussed throughout this thesis), and almost all would benefit from a significant improvement in actuator technology. Devices in this class are the least mature, and therefore this class offers the most opportunity for innovation.

1.3 Low Mechanical Impedance

All three classes (listed above) of machines that interact with humans are required, at times, to permit free motion. This means that the machines must be backdriveable, or have low mechanical impedance. Before this term can be defined, it is necessary to introduce the concept of an interaction port.

> **Definition:** A *port of interaction* (*port, energy port, power port*) for a physical system is an interface through which it can exchange energy with another physical system. The behavior of a port can be characterized by conjugate variables that define power flow (for example, force and velocity in a mechanical system). Ports can be single- or multi-dimensional [14, 108].

Using the concept of an interaction port, properties of the system as it appears at a particular port can be described. One such description is the mechanical impedance.

> **Definition:** The *mechanical driving point impedance* (or simply the *mechanical impedance*) of a physical system at a port (denoted Z) is a dy-

namic operator that determines an output force (torque) time function from an input velocity (angular velocity) time function at the same port. Mechanical admittance at a port (denoted Y) is a dynamic operator that determines an output velocity (angular velocity) time function from a force (torque) time function at the same port. Mathematically, the mechanical impedance $Z = \frac{F}{\dot{x}}$, where F and $\dot{x}$ are the force and velocity, respectively, at the point of interaction. The mechanical admittance $Y = \frac{\dot{x}}{F}$.

Because they are defined with reference to an interaction port, impedance and admittance are generally referred to as "port functions."

Mechanical impedance can also be thought of as a dynamic generalization of stiffness, or as a resistance to motion. If it is desired that a system be backdriveable and permit motion with low resisting force, the system must have low mechanical impedance.

Almost all human-interactive devices have physical interaction as a core function. Some are designed for cooperation or task-sharing with humans, others to take the place of a human in a task. If a robot has high impedance, the human's low impedance cannot significantly change the coupled system dynamics, and the motion of the two is dictated almost entirely by the robot rather than by cooperation. Similarly, a high impedance robot cannot replicate the low-impedance behavior of a human arm. Cooperative tasks between two (or more) dynamic systems often require the systems to have similar impedance, lest one dominate and the other be left useless. Robots that interact with humans must generally have low impedance. For example, if a device is to represent a virtual object, it must provide free motion when removed from the virtual object, so that a contrasting sensation of resistance denotes the object's edges. A telerobotic master typically takes motion of the human operator as an input, and the system moves the slave accordingly, reflecting the forces that the slave encounters back to the master. The human operator must fundamentally be able to provide motion as the system input, and receive force as an output, so the master must provide force without substantial impedance. Extenders and other performance-enhancing devices are designed to actuate at the behest of the human

36

subject, providing assisting forces. As they generally take motion as an input and return force, they must have low impedance. In robotic therapy, a patient with neurological injury attempts to move as directed, and the robot guides his motion to help retrain his neuromuscular system to move properly. The patient must be able to observe his attempts to move and their outcomes, so that he can correct for his errors in subsequent attempts. The robot must move easily in response to even weak inputs. Low mechanical endpoint impedance is thus an essential requirement.

In part because of the endpoint impedance requirements, as well as the force requirements, "high-force haptic" devices such as therapy robots are designed to have dynamics comparable to the dynamics of the limbs with which they interact. Because of this, coupling to the human limb produces coupled system dynamics that are substantially different from the isolated system dynamics, and interaction is, by design, an issue. The impedances are similar, therefore the assumptions that permit analysis as a signal interaction problem cannot be made, and mechanical interaction must be specifically addressed.

1.4   Evaluating Interactive Robots

Like any controlled system, interactive robots must satisfy stability and performance requirements, as well as other practical factors. For an interacting system, each of these has its own particular interpretation. These interpretations are described below, and defined mathematically in subsequent chapters.

1.4.1   Stability

An interactive robot meets its stability requirement if stability is preserved when the robot is coupled to each environment that it encounters in service. If the robot is expected to operate in isolation, free of any environment, it must also be stable in isolation. Not all applications demand isolated stability, however; if the robot is always in contact with some environment, it is not essential that it be stable in isolation. Coupled stability is usually (though not always) a stronger condition than

isolated stability, as it requires stability of a set of coupled physical systems, not just one.

1.4.2 Performance

Performance for an interactive robot consists of several factors. It must be capable of the specified range of motion and forces. Additionally, for most human-interactive applications, it must provide appropriate "feel" at the interface. The relevant endpoint behavior can encompass both static and dynamic effects. Dynamic requirements are analogous to the closed-loop dynamic requirements of conventional robot systems (e.g. rise time, damping ratio, bandwidth, etc.) and can be quantified by the impedance (or admittance) at the point of interaction. Specifically, the robot performs well when its endpoint impedance approaches the target endpoint impedance that provides the desired feel. A performance measure based on impedance is proposed later in this thesis.

1.4.3 Other factors

For human-interactive robots, another important factor is the ergonomic quality of the interface. For many applications, the best interface to the human is one that is low-profile and provides easy attachment. Furthermore, it is generally desirable to minimize overall size. These factors are here noted generally, as the details depend heavily on the particular application. Ergonomics are not directly addressed in this thesis; indirectly, however, minimizing actuator and robot size and mass is an explicit goal that enables improved ergonomic design.

1.5 Evaluating Actuators

Actuators are critical to achieving both robot-level stability and performance. The actuator stability requirement is to ensure that the robot couples stably to its environment. The actuators provide the mechanical energy source and are therefore the

likely source of potential instability. When used with the robot structure, actuators must also provide desired endpoint impedance. Actuator performance is measured by the ability to render the impedance necessary to provide the robot with its target port behavior (the relationship between robot target impedance and actuator target impedance depends on the robot configuration). Because the actuators are a central part of the drivetrain, low robot endpoint impedance is difficult to achieve with high-impedance actuators. In most robot configurations, actuator mass contributes to total robot endpoint impedance (and increases actuator load), so this is also critical to minimize.

An actuator must facilitate stable interaction between robot and environment, provide high forces and low mechanical impedance, be capable of representing desired mechanical impedance, and be reasonably low in mass. Therein lies the key tradeoff that makes this design problem so challenging. Because the actuator mass is critical, the robot-level force requirement becomes a force *density* requirement at the actuator level. Force density is a measure of force capacity per unit actuator mass. Reducing actuator mass usually means reducing actuator force. Alternative strategies, described in the next chapter, usually trade actuator mass for high impedance. Existing technology does not provide a solution with suitably high force density and low impedance for all human-interactive robots. Before such machines can be built, appropriate actuators are needed. This work focuses on improving this tradeoff, exploring the limits of force-based control methods and proposing new hardware configurations to improve performance.

1.6 Actuator for Robotic Therapy

Throughout this work, technology is developed that is versatile in providing actuation and control solutions for a broad set of human-interactive machines, with extension to other machines that interact with their environments. Particular focus is given, however, to the development of a system for robotic therapy; various embodiments of this system are discussed throughout this document to provide examples. The

39

example provides meaningful quantification to the specifications and characteristics of the prototypes and analysis described here. Where the work described specifically addresses the robotic therapy application, an attempt has been made to explain how it generalizes to other applications with other specifications. Here, the application and its requirements are introduced.

The MIT-MANUS robotic therapy aid [62], shown in Figure 1-1, has been used to deliver physical therapy to more than 250 recovering stroke patients over the past decade [1, 40, 42, 75, 76, 114]. MIT-MANUS delivers therapy to the shoulder and elbow via movements in a tabletop plane. Patients play a simple video game that gives them visual cues to move their impaired hand to a particular location. If a patient is unable to move correctly, the robot provides a soft assisting force in the correct direction; the better the patient performs, the less assistance the robot offers. The robot can also be used to provide programmed resistance for more advanced patients to build their strength, and is an effective tool for quantitatively evaluating patient progress. Robotic therapy aids like MIT-MANUS offer several potential advantages, including freeing therapists to work with multiple patients or to give more specialized attention to a single patient, and to provide each patient with more exercise in a typical session (the robot is not subject to fatigue). Robots can reduce the physical intensity of therapists' workload so they can focus more on individual patient needs.

Research studies have consistently shown that patients who receive therapy from MIT-MANUS in addition to their normal regime of physical therapy show a significantly greater reduction in impairment (as determined by clinical scales) than those who receive conventional therapy alone [1, 75, 114]. This result holds equally true for severely-impaired recent stroke victims and for chronic patients, years after stroke [40, 42]. Studies have also consistently shown, however, that improvement is isolated only to the specific muscle groups exercised; the wrist and hand, for instance, show no greater improvement in patients who receive additional shoulder and elbow therapy on the robot.

The results of these studies are clear: robotic therapy is quite effective in reducing impairment, but only in the muscles groups that are specifically targeted. Thus to
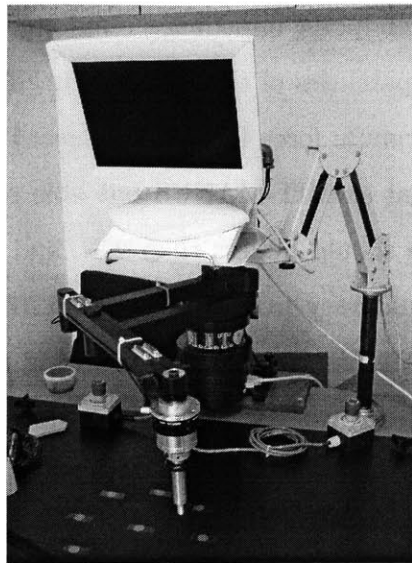
Figure 1-1: MIT-MANUS2 physical therapy robot.

increase the value of the technology it must be extended to exercise more muscle groups.

One of the simplest ways to extend the benefits of MIT-MANUS to a broader set of muscles is to add an additional degree-of-freedom to permit spatial exercise, against gravity as well as in the tabletop plane. This would permit more shoulder exercise than planar therapy, training against gravity, movement training with fewer spatial constraints to more closely approximate free movement, and the re-development of other muscle groups. To create a third degree-of-freedom, a third actuator is needed. One way to do this is to add an actuator to MIT-MANUS, at the point where the patient connects to the robot, to provide vertical motion [16, 15, 74]. This would extend the workspace of the robot from a tabletop plane to a three-dimensional box, permitting the retraining of arm motion against gravity that is essential for basic tasks of daily living such as reaching, feeding, and grooming.

The placement of an extra actuator on the endpoint, rather than redesigning the entire robot, is desirable because it permits modular use of the robot. If planar therapy is desired, the extra actuator can be removed, reducing robot size and inertia. If only therapy against gravity is desired, the actuator can be used separately.

The motion, force, and endpoint impedance requirements for the robotic module are derived from typical capabilities of the human arm. The target range of vertical motion is 14-18". The maximum force is 45 N downward and 65 N up, meeting the same 45 N force requirement as MIT-MANUS but able to also bear arm weight up to 2 kg. The static friction should be less than 2 N, and the inertia less than 2 kg. The endpoint stiffness should vary from 0 to at least 2000 N/m. [20, 15]. Several generations of this device are described in the following chapter, and the hardware and control development described throughout this document are targeted toward meeting these specifications.

1.7 Thesis organization

This thesis seeks to advance the state of the art in interactive machines, particularly in machines that interact with humans, by developing improved hardware and control design methods for modular and scaleable actuators. The objective is to improve the tradeoff between force, mass, and impedance. Prototypes and analysis are targeted to the single degree-of-freedom therapy robot actuator described in the preceding section. Even as this particular application is considered, so is extension to other actuator sizes, requirements, and configurations; where appropriate, detailed descriptions of tradeoffs between design parameters are provided.

This thesis takes a multi-pronged approach to this difficult problem, exploring control strategies, actuator structures, and transmissions. Force-based control methods are explored first. A new measure of stability termed *complementary stability* is proposed that uses limited knowledge of the environment to reduce conservatism, and a design algorithm based on this measure is proposed. This algorithm is validated through simulation and experiments on a single degree-of-freedom robot platform. The simple linear model used for design yields a non-obvious controller that improves both interaction stability and performance versus a state of the art controller. Because force feedback can be insufficient to solve some challenges, changes to the physical actuator structure are considered. The control algorithm is applied as a systematic

means of selecting parameters for dynamics to be placed in series between the actuator and the environment. Several new results are found that help to explain the successes and limitations of series dynamics reported in the literature. Another new finding is that adding inertia in series can have benefits as well. Finally, transmissions are explored as a means of decoupling the actuator weight from the robot endpoint, as well as a way of implementing series dynamics. A hydraulic transmission design for use with a low-impedance source actuator is proposed, and a prototype is constructed and evaluated.

Chapter 2 provides background on the state of the art in interactive machines. In chapter 3, the new stability measure is introduced and a design procedure is proposed. Chapter 4 contains the validation of the methods introduced in chapter 3, in both simulation and hardware implementation. In chapter 5, changes to the physical actuator structure are considered in the context of their ability to improve human-interactive performance, and the tools developed in chapters 3 and 4 are used to determine parameters of dynamics added in series with the actuator. Chapter 6 discusses implementation of series dynamics via a novel fluid transmission intended to improve the interactive performance of conventional actuators. Chapter 7 describes the design details of a prototype transmission, and the test procedures and results used to validate the design. Chapter 8 provides conclusions and a discussion of alternative approaches to the actuation problem.

Chapter 2

State of the Art in Interactive Machines

The force and impedance requirements of interactive machines, as described in the previous chapter, have been addressed with a variety of hardware and control designs. The success of most of these depends on the specific application, and for certain cases, the problem has been elegantly solved. However, a universal solution to the broader high-force, low-impedance, low-mass actuation problem remains elusive. This chapter describes generally the common types of actuators suitable for human-scale force and motion, and evaluates their capabilities in the context of interactive robots. Control strategies for interactive machines are introduced, and the benefits and drawbacks of several control methods for augmenting interactive performance are considered.

2.1 Quantifying actuator and haptic device performance

Ideally a single measure of performance could be used to compare various distinct technologies for actuation. Because a satisfactory actuator for high-force haptic applications must have so many different properties, however, it has proven difficult to define a single measure, or even a small set of measures that comprise a complete

evaluation. An actuator must satisfy force and range-of-motion requirements. It must have low (static and dynamic, linear and nonlinear) friction, inertia (mass of the pieces that move with actuation), and weight (total mass). It must be capable of rendering impedance, both low and high, accurately. To do this, it must have sufficient bandwidth, a control system that responds quickly, and accuracy and repeatability in its output. Finally, it must satisfy other environmental constraints by being suitably packaged, quiet, and clean.

Hollerbach, Hunter, and Ballantyne [64] published a comparative analysis of several actuators on the basis of power density and force density. This work is useful in demonstrating that, for example, hydraulic actuators have higher force density than electromagnetic actuators. They also provide information on other factors of each technology that affect performance, such as limitations in the strain capabilities of polymer actuators. Hayward and Astley [53] provide a long list of performance measures particularly for haptic devices, including the range of motion and force, the inertia, damping, peak acceleration, resolution, precision, and bandwidth. Morrell and Salisbury [83] expand on this list and include both powered (feedback-controlled) and unpowered impedance, force dynamic range, and force fidelity (see [83] for definitions of these terms). They also distinguish between force and position bandwidth, and between force and position precision. Neither Hayward nor Morrell make an attempt to synthesize these many measures into a subset of combined measures; they merely suggest that all measures can be used to characterize a haptic device. Colgate and Brown [25] come closer to this objective by introducing the concept of the "Z-width," a dynamic range for achievable stable impedance. They use measurements of the maximum achievable stiffness and damping under various conditions in a device to quantify this, but do not provide a means of computing the Z-width for the more general case.

As a single measure remains elusive for evaluating actuators, in this work several measures are used. Actuators are evaluated on the basis of their force density. As the specific hardware used for this thesis provides linear motion, this comparison is made under the assumption that all competing technologies must satisfy the range-

of-motion requirements. The other primary consideration is the actuator's ability to accurately render desired impedance, both high and low, specific to an application. In chapter 3, a performance measure based on impedance magnitude for a single specific operating point is introduced. In chapter 5, it is shown that multiple applications of this measure can be combined to evaluate performance over several operating points. However, this measures impedance performance alone, and is independent of the force density or the coupled stability properties of the actuator. An environment-dependent measure of coupled stability is introduced in chapter 3.

In short, the requirements on range-of-motion and environmental factors are treated separately; actuators that cannot meet these requirements are eliminated. The other requirements are combined into the measures of force density, impedance-rendering ability, and coupled stability.

2.2 Existing actuation technologies

The main reason that high-force, low-impedance interactive machines remain so difficult to design and build is the lack of actuators suited to the task. Among available actuators, no single architecture meets all the design requirements or clearly outperforms its competitors. There remains substantial room for innovation in actuator design. Here several common actuation methods, suitable for actuation on the human scale, are reviewed in the context of the requirements of interactive machines.

2.2.1 Direct-drive electromagnetic actuation

Electromagnetic actuators, such as rotary and linear motors and voice coil actuators, can be designed to have extremely low intrinsic impedance. Electromagnetic actuators rely on the force produced by two magnetic fields aligning with each other. In a common arrangement, a permanent magnet (or set of permanent magnets) produces one field, while an electric current passing through a set of wire coils produces a second field. Generally one part is stationary and the other is free to move. The moving part can contain either the permanent magnets (this is the case for most

brushless DC motors) or the coils (as in DC brush motors). The interaction of the magnetic fields produces torque to align the fields of the moving and stationary parts.

Torque increases with the strength of the permanent magnetic field, and is proportional to the coil current. The former is determined by the properties of the magnet material and by the magnetic circuit. This circuit necessarily includes an amount of back-iron to transmit the magnetic flux and a gap, usually filled with air, between the moving and stationary parts. The air gap size is generally limited by manufacturing tolerances and bearing stiffnesses. To increase torque, the current can be increased. However, energy is lost to heat in the coil resistance at the rate of i^2R, where i is the current and R is the resistance of the coils. This heat must be dissipated to avoid damage, and usually limits the torque capability of electromagnetic actuators. Heat can better be dissipated by adding mass to act as a heat sink. The need for back-iron to carry the magnetic flux, and for mass to drain heat, tend to limit the torque capacity of electromagnetic actuators in practical applications. For actuators that rely on magnetic shear forces in the air gap (the most common high-force configuration), a rule of thumb for actuator sizing is that shear force is limited to approximately 10 N per square centimeter of magnet area [112].

To achieve torques close to this limit, a great deal of mass is usually required to act as back-iron and heat sink. Thus electromagnetic actuators tend to have poor force densities. In [64] it is shown that electromagnetic actuators typically are limited to torque densities less than approximately 10 Nm/kg (torque is used rather than force to avoid the dependence on travel).

Electromagnetic actuators can be designed for extremely low intrinsic impedance. Because no contact is allowed between the two magnet surfaces where energy is transduced from the magnetic to the mechanical domain, there is no need for intrinsic mechanical friction. Practically, bearings must be used to maintain the air gap, but the bearing friction is subject only to the limits of engineering the bearing; friction from the force-producing mechanism itself is negligible. In some designs, interaction of the back-iron around which the coils must wrap and the permanent magnetic field can produce cogging: preferred positions of the moving part with respect to the sta-

tionary part. A number of techniques, such as skewed coils and slotless windings are available to minimize cogging. The inertia of the moving part is unavoidable, but is not generally a limiting factor in practical applications. The electromagnetic circuit produces some reflected linear damping (due to the coil resistance) and inertia (due to the coil inductance), but both are well-behaved and generally small. When electromagnetic actuators are current-controlled, they can make extremely high-quality low-impedance torque (or force) sources, with torque (or force) proportional to current input.

With a low-impedance electromechanical actuator and high-quality motion sensors and control, desired impedance can be produced quite effectively (a control strategy for low-impedance hardware is detailed in the section on simple impedance control below). Local "virtual" stiffness is generally limited by sensor resolution, control and power electronics bandwidth, and dynamics between the position sensor and the force-producing surfaces. With careful design, feedback can usually produce actuator stiffness exceeding the requirements for human interaction. Similarly, velocity feedback (from velocity sensors or differentiated position sensors) can be used to add "virtual" damping. As long as the sensors and actuators are collocated, the system can be made quite robustly stable for both contacting and non-contacting applications. The combination of nearly zero intrinsic friction, low intrinsic inertia, and a robust method for adding stiffness and damping makes direct-drive electromagnetics the best of common actuator types at rendering a broad range of accurate impedances.

Closed- and open-chain mechanisms

The principle limitation of direct-drive electromagnetic actuators for interactive robots is their low force density. Thus they are effective when a device can be designed that keeps its actuators stationary. For multiple degrees-of-freedom at the endpoint, this requires a closed-chain kinematic mechanism. An example of a two degree-of-freedom planar closed-chain linkage is shown in plan view in figure 2-1. The two large circles represent two rotary actuators. The angular rotations of the two actuators determine the position of the endpoint, marked by the square, within some limited workspace.

49
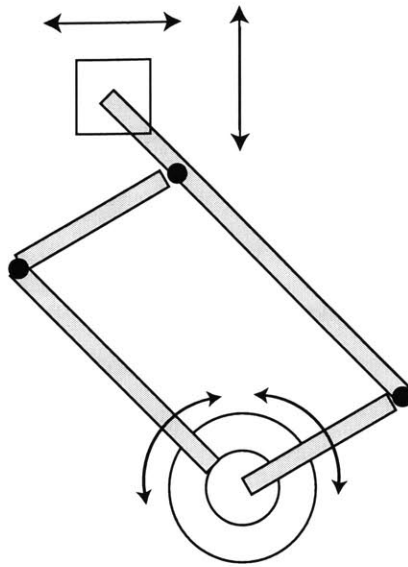
Figure 2-1: A two degree-of-freedom closed-chain kinematic mechanism.

The actuator bodies need not move as the robot is actuated. The majority of the actuator mass can usually be concentrated in the unmoving stator, so that the actuators do not add substantial reflected inertia to the mechanism. Perhaps more importantly, even if the mechanism moves against gravity, the weight of the actuators need not affect the gravitational load at the endpoint.

By contrast, figure 2-2 shows an example of a planar two degree-of-freedom open-chain mechanism. Actuator A is stationary regardless of linkage movement, but actuator B moves with the angular rotation of actuator A. This mechanism has the disadvantage of increased inertia, due to the movement of actuator B, and possibly increased gravitational loads on the structure, on actuator A, and on the endpoint (depending on the orientation with respect to gravity). Because of these higher loads, it is likely that the links must be stronger and therefore more massive. The advantage of such a mechanism, however, is that it can generally have a lower profile and a smaller total package size for the same workspace. This advantage becomes increasingly important as the number of degrees-of-freedom and the complexity of robot endpoint motion increases. For more complex motions, it is difficult if not impossible to design a closed-chain mechanism to complete the task. Many robots that
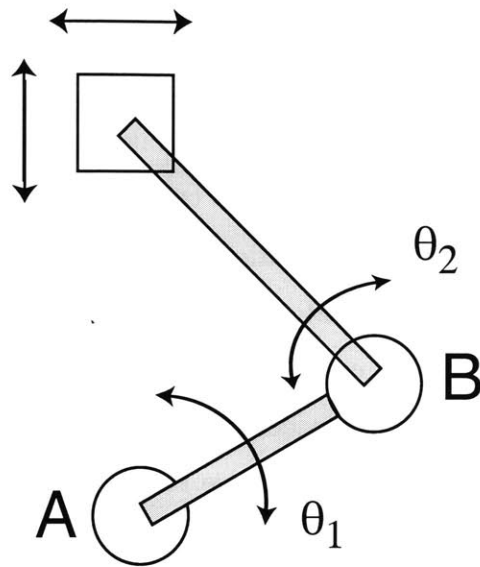
Figure 2-2: A two degree-of-freedom open-chain kinematic mechanism.

are designed to interface with humans must have a low profile to permit the subject to maintain awareness of their environment while working with the machine. As complexity increases, it is almost essential to use mechanisms that carry at least one actuator.

Direct-drive actuators are rarely used as moving actuators on open-chain mechanisms because of their large mass. To reduce their mass sufficiently usually requires reducing their force/torque output to an unacceptably low level. The limitations of this tradeoff usually lead to gearing to improve force density; this strategy is discussed in the following section.

Exoskeletal devices

In addition to posing a problem for open-chain mechanisms, actuator mass and weight is of particular concern for machines that are carried by human subjects, such as exoskeletal devices. Many robots for human interaction are intended for such purposes; examples include active orthoses [11] and rehabilitative devices [117], active leg supports for performance enhancement [90], and powered prosthetic devices [71].

High constant-force loads from actuator weight can be much more severe than

51

inertial loads at modest speeds, and fatigue can be a major concern [68]. This problem rules out direct-drive electromagnetic actuation in all but applications requiring the lowest forces and torques. Despite their outstanding impedance-producing capability, their low force density rules them out almost completely from this class of design.

Because of their force density limitations, direct-drive actuators have generally been used only in particular kinematic configurations where they can be kept stationary, as in MIT-MANUS [62], or for devices that require only very small motions, including magnetically levitating devices [10]. The benefits and drawbacks of direct drive actuation for high-force interactive robots are quantified in the following example:

> **Example 2.1: Linear Motor** A direct-drive linear electromagnetic motor was selected for the vertical therapy robot module application described in chapter 1. A Thrusttube 2504 from Copley was tested for this purpose. A photograph of a commercially produced version of the system (from Interactive Motion Technologies) is shown in figure 2-3. The actuator is an excellent low-impedance force source; this was confirmed by a battery of tests [86]. The intrinsic inertia of the moving forcer is 1.05 kg (with bearings and the handle and force transducer assembly, the total inertia is approximately 2.7 kg). Static and Coulomb friction are between 3 and 4 N, and are quite uniform. Qualitatively, this friction level and this amount of variation in friction feel satisfactory and natural to the human arm. The range of motion is 19 inches, and the continuous force limit is 51 N. The total module mass is 8.1 kg. The device is within or very close to specification except for the total mass, which is more than four times the target mass, and the friction, which is at least 50% above the target, but which may nonetheless suffice. This is a typical result for a direct-drive design; impedance specifications can be met effectively, but the force density is poor.
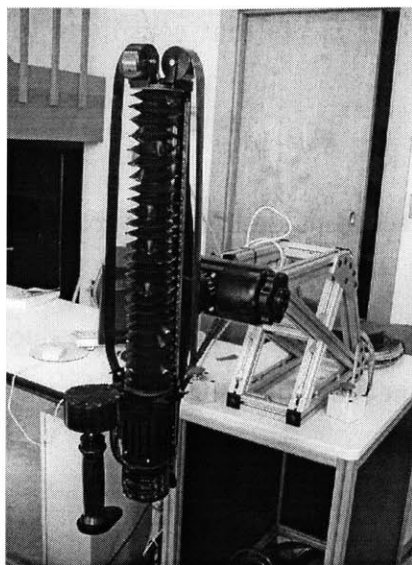
Figure 2-3: Photo of a direct-drive, linear electromagnetic motor-actuated therapy robot module.

2.2.2 Electromagnetic actuators with gears

Almost all traditional robots that are electromagnetically driven use gearing to dramatically improve drivetrain force densities, permitting the use of open-chain mechanisms. Motors are generally sized based on power rather than force requirements, and an appropriate gear ratio is subsequently selected to meet the force specifications. This permits the motors to run at higher speeds, at which most electric motors are dramatically more efficient. This approach can ameliorate the force density limitations of electromagnetic actuators for human interaction, but at a cost, as shown below.

If a rotary motor is modeled as a simple rotary inertia J, with its motion defined by the angle θ and its time derivatives $\dot{\theta}$ and $\ddot{\theta}$, subject to linear damping b, Coulomb friction torque approximated by $C sgn(\dot{\theta})$ (where $sgn(a)$, the signum function, is $+1$ when $a > 0$, 0 when $a = 0$, and -1 when $a < 0$), applied input torque τ_{in}, and load torque τ_l, its motion can be described by:

$$J\ddot{\theta} + b\dot{\theta} + C sgn(\dot{\theta}) = \tau_{in} + \tau_l \tag{2.1}$$

53

Rearranging to show the components that contribute to the torque applied to the load:

$$\tau_l = J\ddot{\theta} + b\dot{\theta} + Csgn(\dot{\theta}) - \tau_{in} \qquad (2.2)$$

If a low-impedance torque source is desired, the goal is to transmit all of the input torque to the output, such that $\tau_l = -\tau_{in}$. The inertia, damping, and Coulomb friction make the source non-ideal. If an ideal gear reduction with ratio N is added between the actuator and the load, the geared displacement θ_g is defined by:

$$\theta_g = \frac{1}{N}\theta \qquad (2.3)$$

and the geared torque at the load τ_g is defined by:

$$\tau_g = N\tau_l \qquad (2.4)$$

Differentiating 2.3, substituting into 2.2, and multiplying by N produces:

$$\tau_g = N^2 J\ddot{\theta}_g + N^2 B\dot{\theta}_g + NCsgn(\dot{\theta}_g) - N\tau_{in} \qquad (2.5)$$

The input torque τ_{in} is scaled by the gear ratio N as desired. Comparison of equations 2.2 and 2.5 reveals that the apparent inertia are scaled by the square of the gear ratio, while the Coulomb friction scales linearly with the gear ratio (as does static friction, because its magnitude does not depend on velocity). As the gear ratio N is increased, so is the endpoint impedance, even if the gearing is implemented ideally, as it is in this example. In fact, most implementations of gear ratios are far from ideal, as discussed below.

In contrast to the interaction problem, high endpoint impedance is not always undesirable in traditional motion-controlled robotics. In fact, an ideal *motion* source has infinite impedance, so high drivetrain impedance can help, so long as it does not adversely impact controllability. For the interaction problem, however, low endpoint impedance is a specific objective. Gearing introduces a painful tradeoff: without it,

the actuator weight may be too large, but with it, the endpoint impedance may be too great. This tradeoff is most evident from a robot-level perspective. Considering again figure 2-2, suppose that actuator A is a direct-drive actuator, while actuator B is geared with ratio N. Assuming that both actuators are selected to satisfy a particular endpoint force requirement, if N is large, actuator B can have low torque capacity and therefore low mass. This reduces the apparent inertia due to changes in θ_1, because the mass of B is small. However, the inertia (and friction) due to changes in θ_2 is increased, as shown in equation 2.5. Conversely, if N is small, actuator B must produce large torques and therefore have high mass, but the impedance properties around θ_2 approach those of a direct-drive actuator as $N \to 1$. But the high mass of actuator B increases the inertia around θ_1.

Understanding the force density / impedance tradeoff in geared electromagnetic actuators is complicated somewhat by the fact that smaller actuators used with large gear ratios have smaller rotor inertia and, in some cases, lower friction than large actuators that produce sufficient torque without gearing. In other words, the apparent endpoint inertia need not increase by a factor of N^2 when a gear ratio is introduced, because the motor inertia J, which produces the apparent inertia $N^2 J$, is smaller for a smaller source motor. To provide a concrete example, a family of practical motors, the Goldline XT line from Kollmorgen, is considered [31]. For a set output torque level of 22 Nm, each of several motors with different frame sizes is assigned a gear ratio that provides the appropriate output torque. From the motor specifications and the gear ratio, the apparent endpoint inertia, friction, and mass are computed for each motor, assuming ideal inertialess, frictionless, and massless gearing. The results are plotted versus gear ratio in figure 2-4.

When apparent impedance properties are adjusted for the changing size of the source motor, the apparent inertia and damping increase steadily with gear ratio, but with less than quadratic dependence. This can be seen in figure 2-4. In general, a larger source actuator with a smaller gear ratio N produces lower endpoint damping and inertia. No obvious trend in reflected static friction torque with N is evident from this data (although most gear implementations are likely to add friction). The
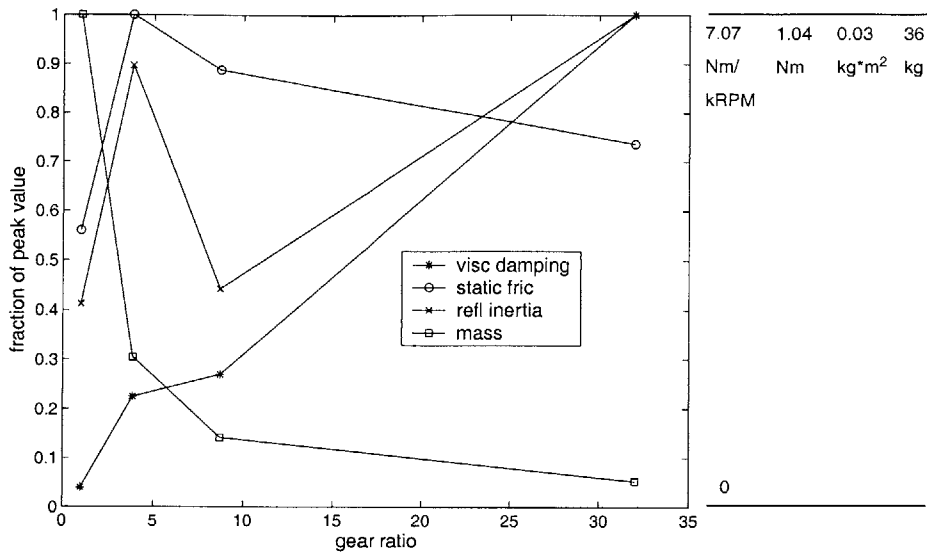
Figure 2-4: Mechanical property tradeoffs when different gear ratios are used, for a fixed output torque of 22 Nm.

relationship between mass and N shows the advantage of ideal gearing, as the actuator mass decreases dramatically with even modest gear ratios. Clearly, the torque/mass ratio improves as the gear ratio increases.

Of course, figure 2-4 is inaccurate in light of the details of any real implementation. The introduction of any gearing with $N \neq 1$ requires a structure with some finite mass. The last three points on the mass curve in figure 2-4 would thus be increased. Similarly, the endpoint inertia must increase from the contribution of the gear system. Most implementations, especially mechanical gears, contribute significant static and Coulomb friction that may be position-dependent (see [5] for a detailed analysis and experimental data from gear friction in a PUMA robot). Other configurations that contribute mechanical advantage can have similar problems (see example 2.2 below). The increase in endpoint impedance resulting from nonlinear gear friction is often the most serious limitation in using gears to increase the force density. Mechanical gears often also introduce backlash, which interferes further with the realistic representation of virtual objects. Alternative implementations of gear reduction (such as cable-pulley systems) may contribute less friction, but introduce other problems, as discussed

56

below.

Several human-interactive devices (e.g. see [122, 73, 117]) use modest mechanical gear reductions to reduce the weight of moving actuators. Others (e.g. [28]) use rotary-to-linear transformations that accomplish the same. The following example quantifies one such system.

> **Example 2.2: Screw-driven vertical module** The original design for the vertical therapy robot application used a rotary motor and a low-friction roller screw to convert to linear motion [15, 16]. The screw provides mechanical advantage analogous to gearing. A photograph of this design is shown in figure 2-5. The force and impedance capabilities of this module were extensively characterized; full results are in [15]. In spite of the extremely low-friction screw design (which theoretically provides pure rolling contact between the nut and screw), position-dependent friction between 10 and 20 N in magnitude is a significant problem. A plot of the screw friction versus position at a variety of speeds is shown in figure 2-6. Friction is strongly periodic with screw rotation, and is also higher when the nut moves along the upper portion of the screw. The variation in friction is particularly disruptive to virtual object representation. The system also has apparent endpoint inertia of almost 6 kg. Although the actuator mass is only approximately 1.1 kg, the total module weight is around 7.7 kg, exceeding the specification by nearly a factor of four. The module is capable of continuous forces up to 139 N and stiffness exceeding 200 kN/m, exceeding the specifications of 65 N and 2000 N/m by 115% and several orders of magnitude, respectively. As compared to the direct-drive module, this version has approximately the same mass but far greater impedance. Its only advantage is a superior force capacity, which is unnecessary for this application.
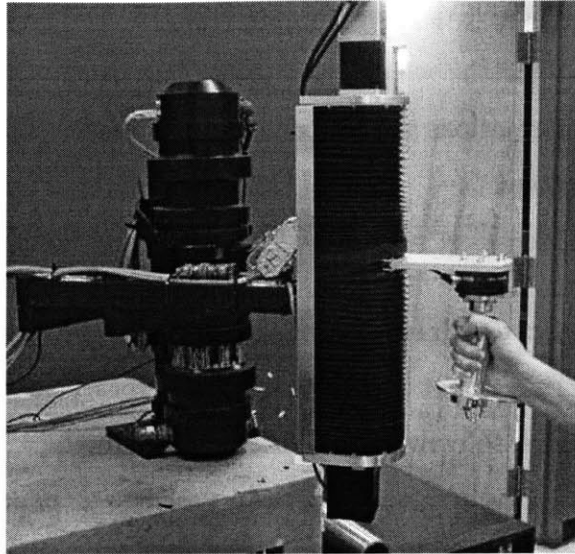
Figure 2-5: Screw-driven vertical robot module, mounted to planar physical therapy robot.
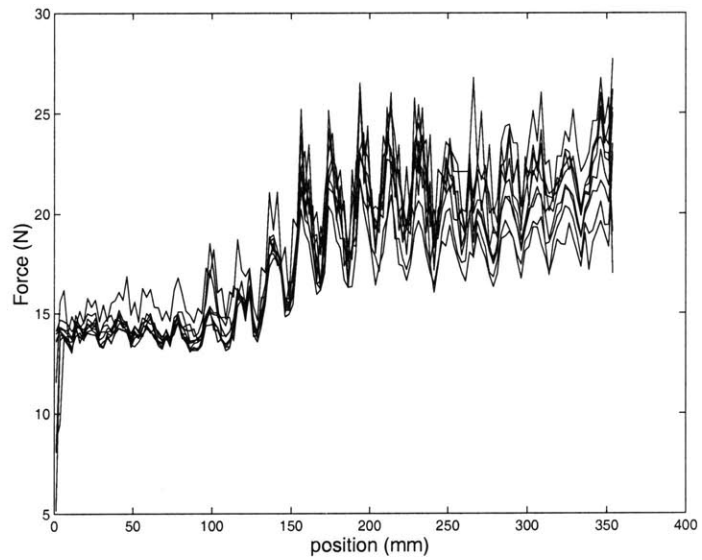
Figure 2-6: Friction force versus position on the screw robot module, at constant speeds ranging from 0.5 to 50 mm/sec.

2.2.3 Electromagnetic actuators with remote transmissions

Remote transmissions promise to transmit the high-quality impedance of electromagnetic actuators to a separate location in a low-weight package. Many common transmissions are made up of tension elements like belts, cables, or tapes, strung over low-friction pulleys or wheels. These systems can have very low endpoint weight, because the source actuator can be physically removed from the point of actuation. Proper design allows the actuator body to be stationary or nearly stationary as the end-effector moves. The only elements that must be at the point of actuation are the tension elements and any wheels or pulleys, bearings, and mounting hardware. This has the potential to produce much higher force densities than any system that requires electromagnetics at the endpoint.

Remote transmissions almost always introduce structural dynamics at resonant frequencies lower than those of the source actuator. This is a nearly unavoidable consequence of transmitting force over a significant distance in a low-profile package. Tension elements, in particular, are often flexible. Resonant modes tend to be numerous, arise in many different forms, and can result in unwanted endpoint vibration. In human-interactive devices, such vibration can be obtrusive and can ruin haptic immersion. Resonant frequencies can be moved with careful design, but can rarely be eliminated.

An added design problem with all remote transmissions, one to which tension elements are particularly vulnerable, is mounting them along the robot linkage. Tension elements must remain loaded in tension and are usually strung along robot arms to ensure this loading, and to avoid obstructing the robot workspace. When many degrees-of-freedom must be actuated with tension elements and many sets of tension elements must be mounted along a robot's arms, this design problem can be challenging. This problem can be solved, sometimes quite elegantly (see e.g. [111]), but this is not a trivial task. Despite their flaws, transmissions using tension elements have been successfully used in several interactive robots [80, 111, 85], and they deserve further study and consideration when such robots are designed.

Alternative transmission designs, such as pushrods, can avoid low-frequency resonance, but can present a substantial design challenge and generally are not as effective at reducing weight or package size.

Passive hydraulic transmissions also fall into this class of actuation technology, and are discussed at length in chapters 6 and 7. Though vulnerable to low-frequency resonance, fluids have some ability to suppress the vibrations that can ruin haptic immersion, while other transmissions like cables can produce or exacerbate these effects. Also, flexible fluid lines can often be "plumbed" more easily than cables to a moving endpoint because they need not keep the same shape to operate properly. Exploration of a specific embodiment of this idea is discussed later in this thesis.

2.2.4 Hydraulics

Hydraulic systems are known for their improved force density versus electromagnetic systems. Hydraulic cylinders and hydraulic motors are mechanisms to convert fluid pressure and flow to mechanical force and motion, and vice versa. Output force or torque is proportional to pressure, and velocity is proportional to the volumetric fluid flow rate. Energy is brought to the actuator in the form of fluid under pressure, potentially in flexible lines, rather than electrical current. Hydraulic systems have higher force densities than electromagnetic systems because the fundamental mechanism for transduction from hydraulic to mechanical energy can have less mass than the mechanism to transduce from electrical to mechanical energy. The higher the fluid pressure, the smaller the effective fluid actuator area can be made without reducing force capacity. However, higher pressures also require stronger structures, which can increase the size and mass at the endpoint. Still, hydraulic systems can have force densities at least an order of magnitude better than direct-drive electromagnetic systems (typical torque densities are around 100 Nm/kg [64]), at comparable power levels.

Figure 2-7 shows a schematic of a typical servo-hydraulic system. A compressor converts electrical power to fluid power, raising the pressure of fluid that it draws from a reservoir. Fluid under pressure is supplied to a servovalve, which regulates the fluid power to an actuator, as commanded by some control system. The actuator,
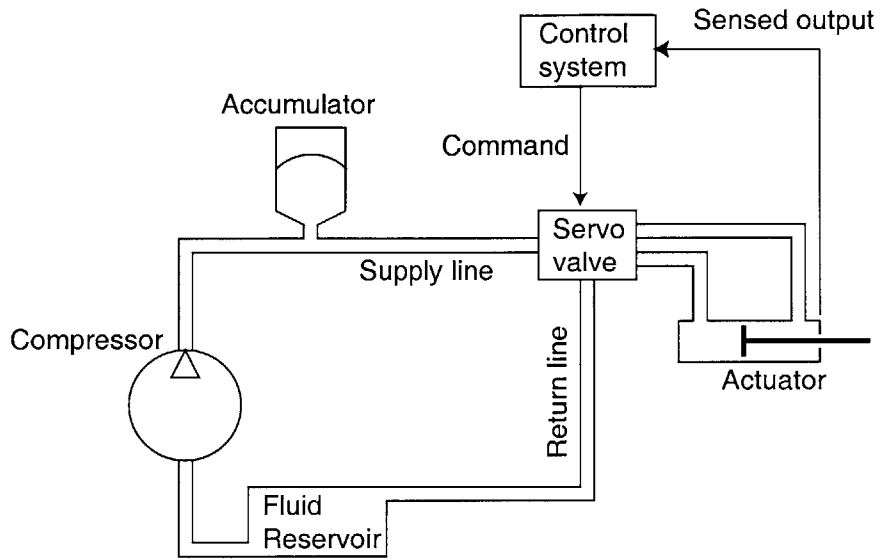
Figure 2-7: A typical servo-hydraulic system.

a hydraulic cylinder or motor, converts from fluid to mechanical power and applies it to the load. Fluid is returned to the reservoir at low pressure. Feedback can be used to measure one or more variables at the actuator and to adjust command to the servovalve through the control system. It is noteworthy that the fluid power delivery requires two stages. The compressor converts from electrical to fluid power, and the servovalve meters it out in response to some command from the control system. This stands in contrast to electromagnetic actuation, where the current (or voltage) is directly modulated, instantaneously adjusting both the total power converted and the command. This distinction indirectly leads to some difficulties in generating desired impedance in hydraulic systems, as shown in the next section.

In principle, servovalves can regulate either flow or pressure. In practice, however, almost all servovalves are fundamentally high-impedance flow regulators. To make a low-impedance pressure source out of such a system, pressure is fed back and flow command is adjusted to meet the target pressure. This arrangement is analogous to using force feedback to make high-impedance mechanical properties disappear, a strategy that is fraught with problems. This approach is analyzed extensively (for mechanical systems) in the following chapters of this thesis.

61

Pressure-regulating servo-hydraulics

Servovalve manufacturers offer valves that are intended for pressure regulation, even without feedback. One example is the Moog Series 15 [81]. This is a two-stage valve; the first stage converts mechanical torque to a differential fluid pressure, given a high pressure source. The second stage is a spool valve configured to act as a purely mechanical pressure amplifier. A schematic of the first stage is depicted in figure 2-8. The rotor is electromagnetically actuated such that torque is proportional to current. Two streams of fluid from the high-pressure source P_s push in opposite directions against the flapper, which is attached to the rotor. Both fluid streams then drip to the return at pressure P_r. The flapper rotates through θ in response to an applied torque to partially restrict one side and raise the fluid pressure in that branch. It settles to equilibrium when the differential pressure between the two sides produces torque equal to that applied by the electromagnetic system. If both flapper nozzles were connected directly to the pure pressure source P_s, movement of the flapper would not affect the upstream pressure P_a or P_b. Thus the constant orifices o_a and o_b are essential to this configuration. They must create a substantial fluid resistance in order for the change in fluid resistance at the flapper to have an effect on P_a and P_b. Unfortunately, any fluid supplied to the load must pass through one of these orifices, increasing the impedance of the source and degrading its quality as a pressure source.

This can be seen if one half of the system in figure 2-8 is modeled as a pair of fluid resistors, as in figure 2-9. The input orifice is modeled as the resistor R_{oa}, and the flapper opening is modeled as the resistor $R_{la}(\theta)$, which introduces the pressure drop between P_a and the return pressure P_r and depends on the rotor angular position θ. Note that if $R_{oa} = 0$, then the output pressure $P_a = P_s$, and the actuated flapper has no effect on output. If a single fixed rotor position $\theta = \theta_0$ is considered, $R_{la} = R_{la}(\theta_0)$. Modeling the output as an impedance $Z_a = \frac{P_a}{Q_a}$, where Q_a is the volumetric output flow rate, the resistor configuration produces:

$$Z_a = \frac{R_{oa}R_{la}}{R_{oa} + R_{la}} \tag{2.6}$$
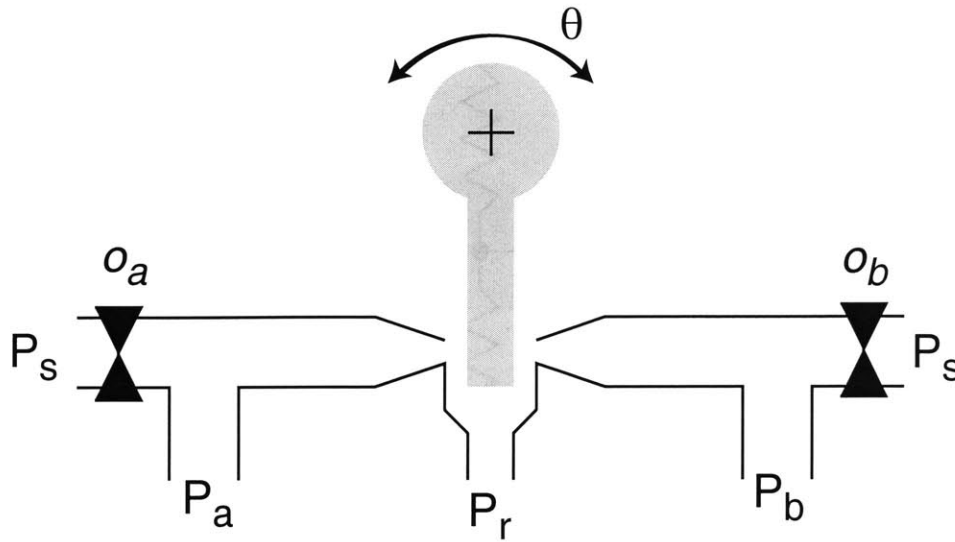
Figure 2-8: Schematic of the first stage of a Moog Series 15 pressure control servovalve.

For this valve to be a pure pressure source, it should have $Z_a = 0$. As noted above, if R_{oa} is too small, then changes in $R_{la}(\theta)$ have little effect on the output pressure P_a. Z_a can also be made small if R_{la} is minimized. However, the total valve output impedance is

$$Z = Z_a + Z_b \qquad (2.7)$$

where Z_b is the port impedance for the other side of the valve, and has the same form as equation 2.6. But $R_{lb}(\theta)$ increases when $R_{la}(\theta)$ decreases, so the only way for the total output impedance to be low is for both flapper resistances to be low. This amounts to allowing substantial leakage through the flapper as compared to the output flow v_a, while also serving to reduce the output pressure. This low output pressure is likely why the Moog design requires a pressure-amplifying second stage. The problems with this amplifier are twofold: first, even if it operates perfectly, it serves to also amplify the impedance of the first stage to the endpoint, as discussed in the section above on gearing (hydraulic gearing is analogous to mechanical gearing); and second, it introduces more small orifices through which the fluid must flow. If these orifices, along with those in the first stage are made large to permit more leakage (by using an open-center spool valve), then substantial flow is required from

63
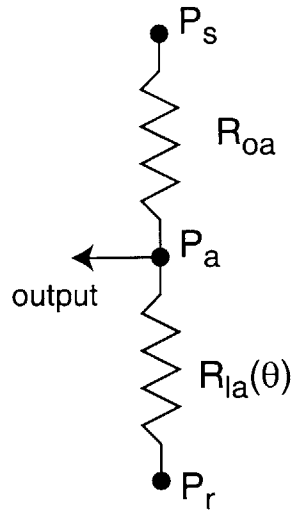
Figure 2-9: Model of one half of the pressure control valve in figure 2-8. The orifice at the flapper is a variable resistor.

the high-pressure source, increasing power consumption and reducing efficiency.

Thus without pressure feedback, the only way to avoid high impedance is with substantial leakage flow, which requires a much more powerful compressor than would be needed if leakage flow is minimized. If a practically sized compressor is used, leakage must be limited to preserve efficiency, and this system is unlikely to act as a good pressure source or produce a high-quality impedance.

It is worth noting that while the servovalve in figure 2-8 requires line resistance to provide pressure output, it does not necessarily require substantial line inertia. In mechanical systems, it has been shown that reducing the inertia via force feedback is the primary obstacle to coupled stability, but that reducing friction may be acceptable if it can be decoupled from inertia reduction (this is discussed at length below). To the author's knowledge, this has not been studied as extensively for hydraulic systems, but it is plausible to speculate that the same may be true for hydraulic systems. If this is the case, then pressure feedback in hydraulic systems using properly designed servovalves may be less problematic for coupled stability than force feedback in mechanical systems. This could aid development of a pressure-controlled hydraulic

actuator that can effectively render desired impedance without compromising coupled stability. This approach is worthy of further study and design effort, and is revisited briefly in chapter 8.

Several other problems, independent of the servovalve, are obstacles to servo-hydraulic systems rendering high-quality impedance. Because pressure is generated separately from the pressure command system, a high-quality constant pressure source is needed (P_s in figure 2-8). Most of the best ways to produce high-pressure use positive-displacement pumps, which inherently are flow sources. An accumulator and a series of pressure relief valves can be used to improve the quality of the pressure source, but maximizing performance is nontrivial. Significantly, the system must be sealed. This is not difficult at the static junctions, but relative motion is required at the actuator. This creates a need for moving seals, which are never perfect and almost always introduce substantial friction. Sealing with low friction is a notoriously difficult problem, and is discussed in chapter 7. Wave transmission effects can be another important factor in hydraulic actuation. When a control loop is closed around long fluid lines, the time delays due to wave transmission can lead to instability [113]. It is better to configure the system such that any long transmission lines fall outside the control loops.

Hydraulic actuation is often used where very high forces are needed, such as in vehicle simulators [103]. Hydraulics have also been used in a dextrous robot arm [66] and in some extenders [70].

2.2.5 Pneumatics

Pneumatic actuation, or air power, is another common method of actuation capable of the force and motion levels required for human interaction. Pneumatic systems use air under pressure to produce mechanical force at a piston or air motor. As in hydraulic systems, the force applied to a piston is equal to the piston area A_p times the differential pressure ΔP across the piston. Piston movement requires air flow in the system.

Pneumatic systems are generally configured similar to hydraulic systems; a com-

pressor converts electric power to air power, and a valve regulates the flow to the system. Pneumatic systems have several advantages over hydraulic systems: their working fluid is plentiful and completely innocuous, so that leakage is not problematic, the fluid itself is not viscous and effectively massless, and does not add friction or mass to the system. Because some leakage is acceptable, extremely low-friction seals can sometimes be used. The primary drawback of pneumatic systems versus hydraulic systems is that air is extremely compressible, and making a stiff system can be difficult. Because the pressure/specific volume curve for air is nonlinear, pneumatic systems become stiffer at higher pressures. To make the system stiff at low force levels, high differential pressures are required. When stiffness is high and motion is swift, this requires the compressor to produce air at a very high pressure and flow rate, which in turn means the compressor must be very powerful. High-powered compressors are noisy, which is another disadvantage in an environment that contains humans. When force levels are relatively low, pneumatics require much higher pressures than hydraulics to achieve high stiffness, meaning that the system components must be stronger and probably more massive. Safety is another concern; because air is highly compressible, it is capable of storing large amounts of energy. If stored energy is released unintentionally, the results can be dangerous.

In general, pneumatic systems provide a soft, comfortable feel to humans. They can typically achieve torque densities of at least 20 Nm/kg, double that of direct-drive electromagnetic actuators, and can achieve even higher levels in certain configurations [64]. They are worth further consideration, especially where high stiffness is unnecessary and safety can be guaranteed. Pneumatic pistons have been used as source actuators for a dextrous hand [54] and a master glove for teleoperation [13], and muscle-type pneumatic actuators have been used in several devices [82, 71].

2.2.6 Interactive actuator design space

Although they each hold some promise, none of the common actuation methods are ideally suited to the human interaction problem, where high force density and the ability to render a high-quality, accurate impedance are both critical. Figure 2-10 pro-

vides a visual representation of the actuator design space for this application, showing existing technology along with desired performance. The horizontal axis is the force density, or the force/mass ratio. The vertical axis represents the ability to represent a broad range of impedance with high quality. To qualify as "excellent" (in figure 2-10) at representing impedance, an actuator must have low intrinsic impedance, be capable of stiff impedance with motion feedback, and display forces with high fidelity. Electromagnetic actuators are alone among the major classes in their ability to meet both of these requirements, but they have extremely poor force density; thus they are placed in the upper left-hand corner of the plot. A cluster of other actuation technologies have much better force densities, but each has a flaw in their ability to render impedance (as described above); this places them in the lower right-hand corner. Geared electromagnetic actuators all lie along a curve on this plot, depending on the gear ratio. Low gear ratios place the system close to direct-drive actuators; very high ratios dramatically increase force density, but increase intrinsic impedance, placing such systems in the lower right-hand corner. The arrow indicates that geared systems with various gear ratios fall anywhere between these two extremes. The upper right-hand corner represents the desired region of performance. None of the actuation technologies discussed here satisfy these requirements intrinsically.

Several actuation technologies, namely highly geared mechanical actuators and hydraulic actuators, are limited mainly by their high intrinsic impedance. If the appearance of this impedance at the endpoint can be artificially reduced, these methods might meet the requirements. One way to reduce the apparent impedance is to measure interaction force and use the actuators to move the system in the direction of applied force, as if it were not there. To understand the viability and limits of this approach, it is necessary to explore the control techniques for regulating physical interaction between a robot and another physical system, its environment.
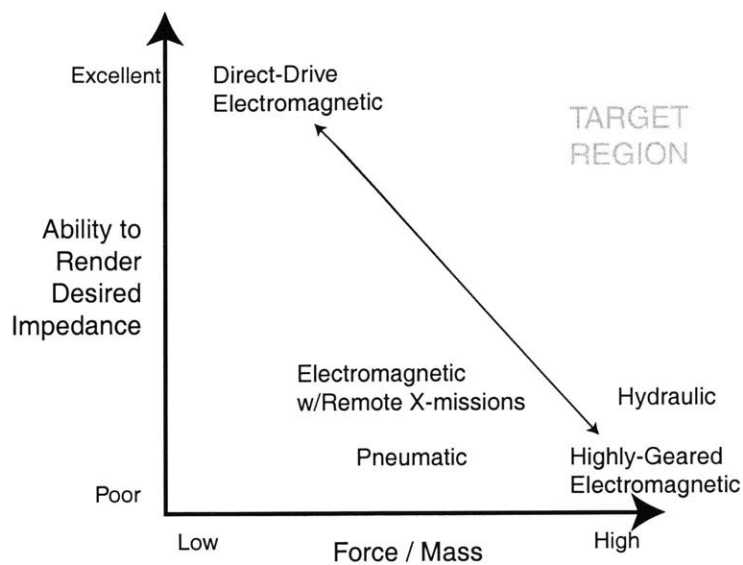
Figure 2-10: Impedance-rendering ability versus force density for various types of actuators. For an actuator to render target impedance effectively, it must be capable of both low and high tunable endpoint impedance.

2.3 Control for interactive applications

To this point, it has been stated that low intrinsic impedance is critical to implementing high-quality impedance behavior across a broad range of impedances. The reason for this, and strategies for regulating interaction under these circumstances, are described in the rest of this chapter [1]. This section also discusses how to make a robot behave as desired, to exhibit the target feel. Finally, existing strategies for using force measurement to reduce the apparent impedance are discussed, and the benefits and limitations of these approaches are analyzed.

2.3.1 Effect of interaction on performance and stability

When interaction occurs, the dynamic properties of the environment are important. If we attempt to control motion or force, interaction affects the controlled variable, introducing error upon which the controller must act. Errors clearly mar performance but even more important, though perhaps less obvious, stability may also be compro-

---

[1]Parts of this section are published in [61].

mised. That is, a system that is stable in isolation can become unstable when coupled to an environment that is itself stable. Such instabilities, known as coupled or contact instabilities, appear even in simple systems with basic controllers contacting simple environments. To illustrate this key point consider the following examples.

Example 2.3: Integral motion control This example is based on a common design for a robot motion controller. One of the simplest models of a robot is depicted in figure 2-11. A single mass m represents the robot's inertial properties. It is subject to actuator forces F_a and environmental forces F_e. A single damper b connected to ground represents frictional losses. The Laplace-transformed equation of motion for this simple model is as follows:

$$(ms^2 + bs)X = F_a + F_e \tag{2.8}$$

where X is the Laplace transform of the mass position, x. A proportional-integral motion controller is applied:

$$F_a = K_p(R - X) + \frac{K_i}{s}(R - X) \tag{2.9}$$

where R is the Laplace transform of the reference position and K_p and K_i are proportional and integral gains, respectively. In isolation, $F_e = 0$ and the closed-loop transfer function is:

$$\frac{X}{R} = \frac{K_p s + K_i}{ms^3 + bs^2 + K_p s + K_i} \tag{2.10}$$

From the Routh-Hurwitz stability criterion, a condition for *isolated* stability is the following upper bound on the integral gain:

$$K_i < \frac{bK_p}{m} \tag{2.11}$$

However, this condition is not sufficient to ensure that the robot will remain stable when it interacts with objects in its environment. Even the

69
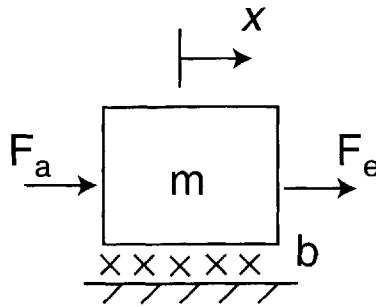
Figure 2-11: Single-mass interactive robot model with linear damping.

simple act of grasping an object may be destabilizing. If the system is not isolated, and instead is coupled to a mass m_e, this is equivalent to increasing the mass from m to $(m + m_e)$. Hence a condition for stability of the coupled system is:

$$K_i < \frac{bK_p}{(m + m_e)} \qquad (2.12)$$

Any controller that satisfies equation 2.11 results in a stable isolated system. Prudent controller design would use a lower integral gain than the marginal value, providing robustness to uncertainty about system parameters and perhaps improved performance. However, for any fixed controller gains, coupling the robot to a sufficiently large m_e will violate the condition for stability. Interaction with a large enough mass can always destabilize a system that includes integral-action motion control, even if that system is stable in isolation.

This difficulty may be obscured by the fact that, at present, most robots have inertia far exceeding that of the payloads they may carry and the total mass $(m + m_e)$ is not much greater than the mass of the robot alone m and the bounds of equations 2.11 and 2.12 are similar. However, in some applications, e.g., in space or under water, a robot need support little or none of a payload's weight and objects of extremely large mass may be manipulated. In these situations the vulnerability of integral-action motion control to coupled instability may become an important consideration.

70

Example 2.4: Force feedback control This example illustrates one of the common difficulties of force control. A robot with a single structural vibration mode is represented by the model shown in figure 2-12. Two masses, m_1 and m_2 are connected by a spring of stiffness k. Frictional losses are represented by dampers b_1 and b_2 connected from the masses to ground and damper b_3 in parallel with the spring. One mass is driven by the actuator force F_a and the other is subject to F_e, an interaction force with the environment.

A proportional-derivative (PD) controller acting on the error between the position of the mass at the actuator x_1 and a reference r is applied to control motion. A proportional controller acting on force fed back from the environment is applied to control force and improve interactive behavior. The control law is as follows:

$$F_a = K(r - x_1) + B(\dot{r} - \dot{x}_1) + K_f(F_e + K(r - x_1) + B(\dot{r} - \dot{x}_1)) \quad (2.13)$$

K and B are related to the proportional and derivative motion feedback gains, respectively, and K_f is the proportional force feedback gain (the benefits of this controller are further discussed below). When the system is isolated from any environment, $F_e = 0$ and equation 2.13 reduces to PD motion control. Using the Routh-Hurwitz criterion the isolated system's closed-loop characteristic polynomial can be shown to have all its poles in the open left half plane, providing at worst marginal stability for arbitrary non-negative real parameters and controller gains.

If the system is not isolated and instead connected to a spring to ground, such that $F_e = -k_e x_2$, the closed loop characteristic polynomial is changed. It is now easy to find parameters such that this polynomial has right-half-plane roots. For example, if $m_1 = m_2 = 10$ kg, $b_1 = b_2 = b_3 = 1$ Ns/m, $k = 100$ N/m, $K = 10$, $B = 1$, $K_f = 10$, $k_e = 100$ N/m, the closed loop poles are at $-2.78 \pm 5.51i$ and $2.03 \pm 5.53i$, where the latter pair are

71
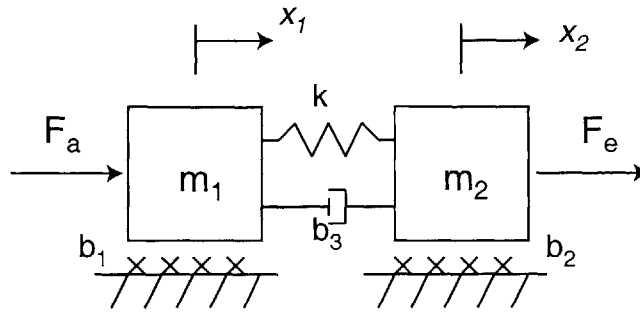
Figure 2-12: A simple robot model with a single structural resonant mode.

unstable. Once again, a system that is stable when isolated (corresponding to $k_e = 0$) is driven to instability when coupled to an object in its environment, in this case, a sufficiently stiff spring ($k_e = 100$).

Structural vibrations are a common feature of robot dynamics and present a substantial challenge to controller design. Though simple, the PD motion controller in this example has the merit that it is not destabilized by structural vibration because the actuator and assumed motion sensor are collocated. However, the addition of a force feedback loop renders the robot control system vulnerable to coupled instability. In part this is because the actuator and force sensor are not collocated. Given the difficulty of designing a robot with no significant dynamics interposed between its actuators and the points at which it contacts its environment (though see [110, 106]) the simple model in this example represents a common situation in robot interaction control and will be a focus of analysis throughout this thesis.

These two examples show that to ensure stability when interacting even with simple environments it is not sufficient to design for stability of the manipulator in isolation. The environment's dynamics must also be considered - dynamics that are generally not known exactly. Furthermore, in many applications a robot must be capable of stable interaction with a variety of environments. For example, an assembly robot might pick up a component, move it across free space, bring it into contact with a kinematic constraint used to guide placement of the component, move it along the constraint, release the component and return to get another. In this case,

at different times in the process the robot must remain stable when moving unloaded in free space, moving to transport the component, and moving to comply with the kinematic constraint. Each of these three contexts poses a different stability challenge which illustrates one key reason why, even if the environment can be closely modeled, there are strong benefits to designing a controller that is insensitive to ignorance of its properties. If, for example, a different controller is required for each of the three contexts, the system must decide which controller to use at each instant and manage the transition between them. Usually the easiest way to identify the environment is to interact with it; thus interaction might be required before the appropriate controller is in place. Stable interaction is needed to identify the environment, but stability cannot be guaranteed without a well-characterized environment. A single controller that can perform satisfactorily within all expected contexts without compromising stability would have significant advantages.

2.3.2 Interaction as disturbance rejection or modeling uncertainty

Control theory offers several tools to endow controllers with the ability to deal with unknown or poorly characterized interference [79]. Using a disturbance rejection approach, the environment's dynamics could be included as disturbance forces. The success of this approach depends on bounding the disturbance forces but for many interactive applications the environmental forces may equal or exceed the robot's nominal capacity; for example a kinematic constraint can introduce arbitrarily large forces depending on the robot's own behavior. Furthermore, environmental forces generally depend on the robot's state, and traditionally disturbances are assumed to be state-independent. Thus treating interaction as a disturbance rejection problem does not seem promising.

Modeling the environment as an uncertain part of the robot and using robust control tools [79, 104] to guarantee stability is another reasonable approach. Interacting with an environment effectively changes the robot plant by adding some combination

73

of elastic, dissipative and inertial properties, perhaps including kinematic constraints. If the effect of interaction is only to alter the parameters of the robot model (e.g., by adding to the endpoint mass) a robust control approach may succeed, though robustifying the system to a large range of environment parameters might require an unacceptable sacrifice of performance. However, interaction may also change the *structure* of the model. For example, it may reduce the order of the system (e.g., moving from free motion to contact with a kinematic constraint reduces the degrees of freedom of the parts contacting the constraint) or increase it (e.g., contact with an elastic object adds a new mode of behavior due to interaction between the robot inertia and the object elasticity). If interaction changes the model structure, the applicability of a robust control approach is unclear. For example, the environment forces and motions may be of the same magnitude and in the same frequency range as the known robot dynamics. This is particularly likely in the case of human-interactive robotics, as robots are frequently designed to have dynamics comparable to those of their expected environments. As robust control methods commonly assume that uncertain dynamics lie outside the frequency range of interest, treating interaction as a robustness problem does not seem promising.

2.3.3 Regulating dynamic behavior

The general approach here termed "interaction control" refers to regulation of the robot's dynamic behavior at its ports of interaction with the environment. An interaction port is a place at which energy may be exchanged with the environment. Interaction control involves specifying a dynamic relationship between motion and force at the port, and implementing a control law that attempts to minimize deviation from this relationship.

Mechanical impedance and admittance

The most common forms of interaction control regulate the manipulator's impedance or admittance [57, 60]. Assuming force is analogous to voltage and velocity is analo-

74
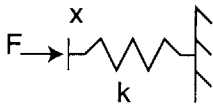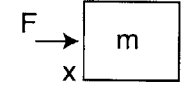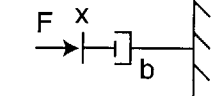
	Impedance $F/\dot{x}$	Admittance $\dot{x}/F$
$F \rightarrow \mathord{\text{/\kern-2pt\backslash}}$ spring k	$\dfrac{k}{s}$	$\dfrac{s}{k}$
$F \rightarrow$ mass m, x	$m\,s$	$\dfrac{1}{m\,s}$
$F \rightarrow$ damper b, x	b	$\dfrac{1}{b}$

Figure 2-13: Examples of mechanical impedance and admittance for simple linear mechanical elements.

gous to current, mechanical impedance is analogous to electrical impedance, characterized by conjugate variables that define power flow, and defined in chapter 1.

If the system is linear and invertible, admittance is the inverse of impedance, and both can be represented in the Laplace domain by transfer functions, $Z(s)$ or $Y(s)$. Figure 2-13 gives example admittances and impedances for common mechanical elements. For a linear system impedance can be derived from the Laplace-transformed dynamic equations by solving for the appropriate variables.

While most often introduced for linear systems, impedance and admittance generalize to nonlinear systems. The port function analysis in this thesis is restricted to the linear case; additional information on nonlinear impedance can be found in [57].

A key point to note is that, unlike motion or force, the dynamic port behavior is exclusively a property of the robot system, independent of the environment at that port. This is what gives regulating impedance or admittance its appeal. Motion and force depend on both robot and environment as they meet at an interaction port, and cannot be described or predicted in the absence of a complete characterization of both systems. Indeed, this is the principal difficulty, as illustrated in the examples above, in regulating either of these quantities. Impedance, on the other hand, can

(in theory) be held constant regardless of the environment; impedance defines the relationship between the power variables, and does not by itself determine either.

Impedance serves to completely describe how the manipulator will interact with a variety of environments. In principle, if arbitrary impedance can be achieved, arbitrary behavior can be achieved; it remains only to sculpt the impedance to yield the desired behavior. Of course, as with all controller designs, the goal of achieving a desired impedance is an ideal; in practice, it can be difficult. The problem of achieving a desired impedance (or admittance) is central to the study of interaction control, and is discussed in subsequent sections of this chapter and throughout this thesis.

Port behavior and transfer functions

Port impedance and admittance are just two of the possible representations of a linear system's response, and it is important to highlight the distinction between these expressions and general input/output transfer functions. By way of example, consider again the system in figure 2-12. This system is an example of a "2-port" because it has two power interfaces, one characterized by F_a and x_1, the other by F_e and x_2. If such an element is part of a robot, one side is generally connected to an actuator and the other to a downstream portion of the robot or directly to the environment. Any mechanical 2-port has four transfer functions relating motion to force. Two of them, $\frac{F_a}{x_2}(s)$ and $\frac{F_e}{x_1}(s)$ (or their inverses), are input/output transfer functions, and define the force produced by motion at the opposite port, assuming certain boundary conditions to determine the other power variables. This type of transfer function is used in traditional block-diagram analysis; if the left port is driven by a controllable force and the right uncoupled, then $\frac{\dot{x}_2}{F_a}$ describes the motion of the right port as a result of actuator force, as shown in figure 2-14. The other two transfer functions, $\frac{F_a}{x_1}(s)$ and $\frac{F_e}{x_2}(s)$, each represent the impedance at a port (and their inverses the corresponding admittance), depending on the boundary conditions at the other port.

The principal distinction between these two types of transfer functions is in their connection. If we have two systems like that shown in figure 2-12, their input/output
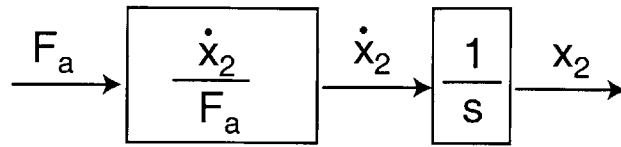
Figure 2-14: A block-diagram representation of a forward-path transfer function for the system in figure 2-12.

transfer functions cannot be properly connected in cascade, because the dynamics of the second system affect the input-output relationship of the first. At best, cascade combination of input/output transfer functions may be used as an approximation if the two systems satisfy certain conditions, specifically that the impedance (or admittance) at the common interaction port are maximally mismatched to ensure no significant loading of one system by the other. This is a common assumption when a controller transfer function is connected in cascade to a plant input/output transfer function, with the assumption that the plant dynamics do not affect the input/output characteristics of the controller. Port functions, on the other hand, can be connected without assumption of their dynamic characteristics. Connection of port functions is discussed below.

Dynamic behavior and "feel"

As mentioned in chapter 1, mechanical impedance and admittance provide representations of a robot's feel, which is a key performance objective for human interaction. Regulating port behavior, unlike regulating motion or force directly, means regulating a quantity that directly expresses performance. Rather than attempt to overwhelm the consequences of interaction by building in robustness, interaction control aims to regulate the interaction itself by controlling the robot's dynamic behavior at the places where it interacts with its environment. Hence robot performance can be judged solely in terms of the appearance of the robot's dynamics to the human, as measured by its port functions, rather than on the basis of motion or force variables

77

$$S_1 \longleftarrow 1 \longrightarrow S_2$$
$$\dot{x}$$
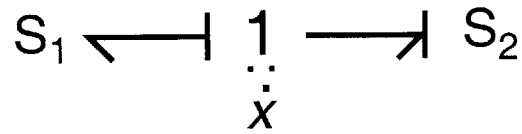
Figure 2-15: A bond graph representation of two physical systems represented by their port functions, S_1 and S_2.

that depend on both the robot and environment.

2.3.4 Analyzing coupled systems

To analyze systems that are coupled by interaction ports, it is helpful to consider the causality of their interaction behavior, that is, the choice of input and output for each of the systems. The appropriate causality for each system is constrained by the nature of the connection between systems, and in turn affects the mathematical representation of the function that describes the system. Power-based network modeling approaches (such as bond graphs [14]) are useful, though not essential, to understand this important topic.

Causal analysis of interaction port connection

Figure 2-15 uses bond graph notation to depict direct connection of two mechanical systems, S_1 and S_2, such that their interaction ports have the same velocity (the most common connection in robotic systems), denoted by the 1 in the figure. If we choose to describe the interaction dynamics of one system as an impedance, causal analysis dictates that the other system must have admittance causality. The causal stroke indicates that S_1 is in impedance causality; in other words, S_1 takes motion as an input and produces force output, and S_2 takes force as an input and produces motion output. If we use the common sign convention that power flow is positive into each system (denoted by the half arrows in the figure) this is equivalent to a negative feedback connection of the two power ports, as shown in figure 2-16.

More complicated connections are also possible. Consider three mechanical sys-
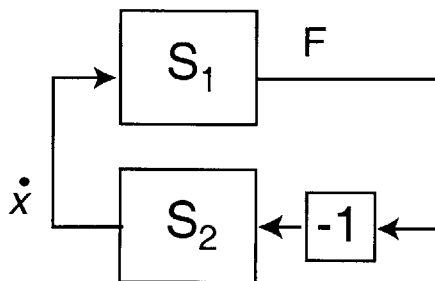
Figure 2-16: A block diagram representation equivalent to the bond graph representation in figure 2-15.
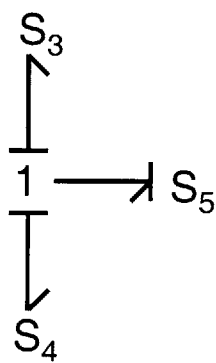
Figure 2-17: A bond graph representations of three systems coupled with a common motion.

tems connected such that their common interaction ports have the same velocity (e.g., consider two robots S_3 and S_4 pushing on the same workpiece S_5). Figure 2-17 shows a bond graph representation. If we choose to describe the workpiece dynamics as an admittance (and this may be the only option if the workpiece includes a kinematic constraint) then causal analysis dictates that the interaction dynamics of the two robots must be represented as impedances. Once again, this connection may be represented by a feedback network. S_3 and S_4 are connected in parallel with each other and with the workpiece S_5 in a negative feedback loop, as shown in figure 2-18. Note that S_3 and S_4 can be represented by a single equivalent impedance, and the system looks like that in figures 2-15 and 2-16. S_3 and S_4 need not represent distinct pieces of hardware, and in fact might represent superimposed control algorithms acting on a single robot [3, 57].

79
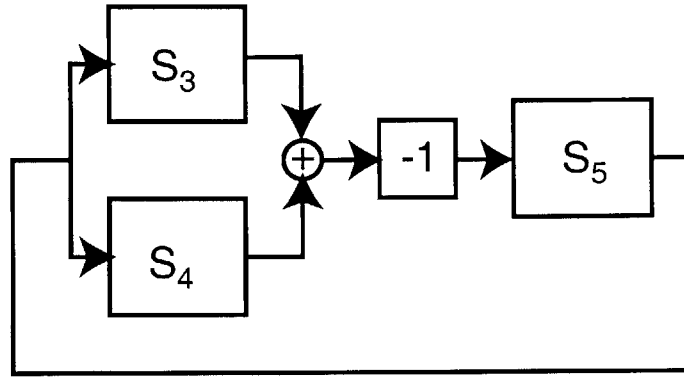
Figure 2-18: A block diagram representation equivalent to the bond graph representation in figure 2-17.

Impedance vs. admittance regulation

Causal analysis provides insight into the important question whether it is better to regulate impedance or admittance. For most robotics applications, the environment consists primarily of movable objects, most simply represented by their inertia, and surfaces or other mechanical structures that kinematically constrain their motion. The interaction dynamics in both cases may be represented as an admittance. An unrestrained inertia determines acceleration in response to applied force, yielding a proper admittance transfer function (see figure 2-13). A kinematic constraint imposes zero motion in one or more directions regardless of applied force; it does not allow representation as an impedance. Inertias prefer admittance causality; kinematic constraints require it. Because the environment has the properties best represented as an admittance, the ideal robot behavior is an impedance, which can be thought of as a dynamic generalization of a spring, returning force in response to applied displacement. Initially assuming that arbitrary port impedance could be achieved, Hogan argued for this approach [57].

However, most robots consist of links of relatively large mass driven by actuators. As described above, if the actuators are geared to amplify motor torque, the total inertia apparent at the end-effector is increased by the reflected inertia of the motor; indeed for high gear ratios this can dwarf the mass of the links. These inertial proper-

80

ties are difficult to overcome (see section below on Natural Admittance Control), and tend to dominate the robot's response. Thus it is difficult to make a robot behave as a spring (or impedance) and it is usually more feasible to make it behave as primarily a mass (or admittance), an argument for admittance control [87]. In other words, impedance behavior is usually ideal, but admittance behavior is often more easily implemented in real hardware, which itself prefers admittance causality because of its inertial nature. The choice of particular controller structure for an application must be based on the anticipated structure of environment and manipulator, as well as the way that they are coupled.

Coupled system stability analysis

If S_1 and S_2 are a linear time-invariant impedance and admittance, respectively, figures 2-15 and 2-16 depict the junction of the two systems via a power-continuous coupling; in other words the coupling is lossless and does not require power input to be implemented; all power out of one system goes into the other(s). Colgate has demonstrated the use of classical feedback analysis tools to evaluate the stability of this coupled system, using only the impedance and admittance transfer functions [23, 27]. By interpreting the interaction as a unity negative feedback as shown in figure 2-16, Bode and/or Nyquist frequency response analysis can be applied. The "open-loop" transfer function in this case is the product of admittance and impedance $S_1 S_2$. Note that, unlike the typical approximate cascade representation of transfer functions, this product is exact, independent of whether or how the two systems exchange power or "load" each other. The Nyquist stability criterion ensures stability of the coupled (closed-loop) system if the net number of clockwise encirclements of the -1 point (where magnitude = 1 and phase angle = $\pm 180°$) by the Nyquist contour of $S_1(j\omega)S_2(j\omega)$ plus the number of poles of $S_1(j\omega)S_2(j\omega)$ in the right half-plane is zero. The slightly weaker Bode condition ensures stability if the magnitude $|S_1(j\omega)S_2(j\omega)| < 1$ at any point where $\angle(S_1(j\omega)S_2(j\omega)) = \angle(S_1(j\omega)) + \angle(S_2(j\omega)) = \pm 180°$.

Passivity and coupled stability

This stability analysis requires an accurate representation of the environment as well as the manipulator, something that is difficult to obtain, and undesirable to rely upon. The principles of this analysis, however, suggest an approach to guarantee stability if the environment has certain properties-particularly if the environment is passive.

There are a number of definitions of a system with passive port impedance, all of which quantify the notion that the system cannot, for any time period, output more energy at its port of interaction than has in total been put into the same port for all time. Linear and nonlinear systems can have passive interaction ports; here only the linear time-invariant case is analyzed, and the definition from Colgate [23] is used. For nonlinear extensions, see [121, 123].

> **Definition:** A system defined by the linear 1-port impedance function
> $Z(s)$ is passive iff:
>
> 1. $Z(s)$ has no poles in the right half plane.
>
> 2. Any imaginary poles of $Z(s)$ are simple, and have positive real residues.
>
> 3. $Re(Z(j\omega)) \geq 0$.

These requirements ensure that $Z(s)$ is a positive real function, and lead to the following interesting and useful extensions:

1. If $Z(s)$ is positive real, so is its inverse, the admittance function $Y(s) = Z^{-1}(s)$, and $Y(s)$ has the same properties.

2. If equality is restricted from condition 3, the system is dissipative and is called strictly passive.

   - The Nyquist contours of $Z(s)$ and $Y(s)$ each lie wholly within the closed right half-plane.

   - $Z(s)$ and $Y(s)$ each have phase in the closed interval between $-90°$ and $+90°$.

3. If condition 3 is met in equality the system is passive (but not strictly passive).

- The Nyquist contours of $Z(s)$ and $Y(s)$ each lie wholly within the open right half-plane.

- $Z(s)$ and $Y(s)$ each have phase in the open interval between $-90°$ and $+90°$.

Note that pure masses and springs, as shown in figure 2-13, are passive but not strictly passive. Pure dampers are strictly passive, as they have zero phase. Furthermore, any collection of passive elements (springs, masses, dampers, constraints) assembled with any combination of power-continuous connections is also passive. Passive systems comprise a broad and useful set, including all combinations of mechanical elements that either store or dissipate energy without generating any-even when they are nonlinear. Passivity and related concepts have proven useful for other control system applications, including robust and adaptive control [79].

Colgate has shown that the requirement for a manipulator to interact stably with any passive environment is that the manipulator itself be passive [23, 27]. The proof is described here intuitively and informally. If two passive systems are coupled in the power-continuous way described above and illustrated in figure 2-15, the phase property of passive systems constrains the total phase of the "open-loop" transfer function to between $-180°$ and $+180°$ (the phase of each of the two port functions is between $-90°$ and $+90°$, and the two are summed). Because the phase never crosses these bounds, the coupled system is never unstable and is at worst marginally stable. This holds true regardless of the magnitudes of the port functions of both systems. In the Nyquist plane, as shown in figure 2-19, since the total phase never exceeds $180°$, the contour cannot cross the negative real axis, and therefore can never encircle -1, regardless of its magnitude. This result shows that if a manipulator can be made to behave with passive driving point impedance, coupled stability is ensured with all passive environments, provided the coupling obeys the constraints applied above. The magnitude of the port functions is irrelevant; if passivity (and therefore the phase
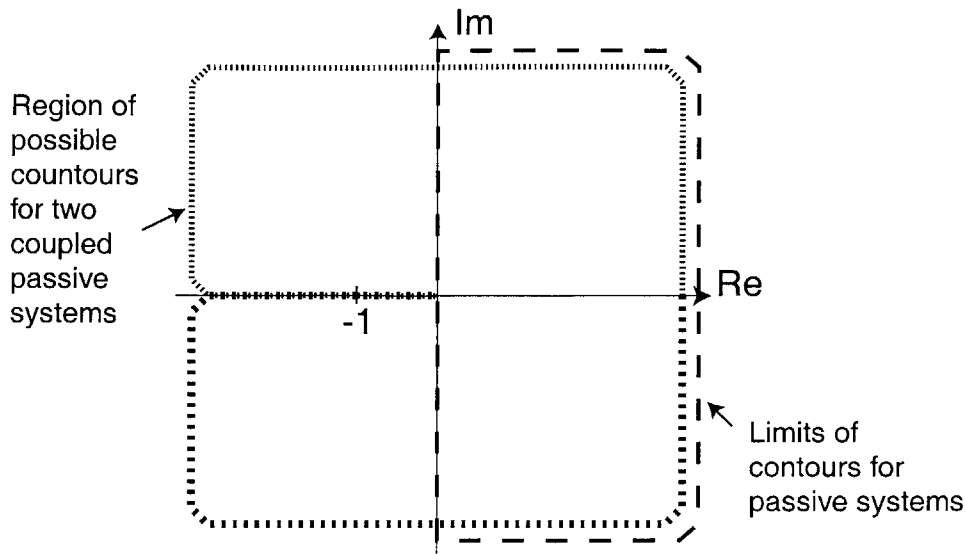
Figure 2-19: Limiting regions of the Nyquist plane for passive systems and for two coupled passive systems.

constraint) is satisfied, the coupled system is stable. The indifference to magnitude also means that requiring a robot to exhibit passive interaction behavior need not compromise performance; in principle the system may be infinitely stiff or infinitely compliant.

It is worthy of mention that if either of the two systems is strictly passive, the total phase is strictly less than ±180°, and the coupled system is asymptotically stable. Different types of coupling can produce slightly different results. If the act of coupling requires energy to be stored, or if contact is through a mechanism such as sliding friction, local asymptotic stability may not be guaranteed. However, the coupled system energy remains bounded. This result follows from the fact that neither system can generate energy or supply it continuously; the two can only pass it back and forth and the total energy can never grow.

Example 2.5: Nonpassive force feedback system Because examples 2.3 and 2.4 both show systems that can be driven unstable when interacting with passive elements, both systems must be non-passive. This is in fact true; for the case of the system in example 2.4, for instance, figure

84
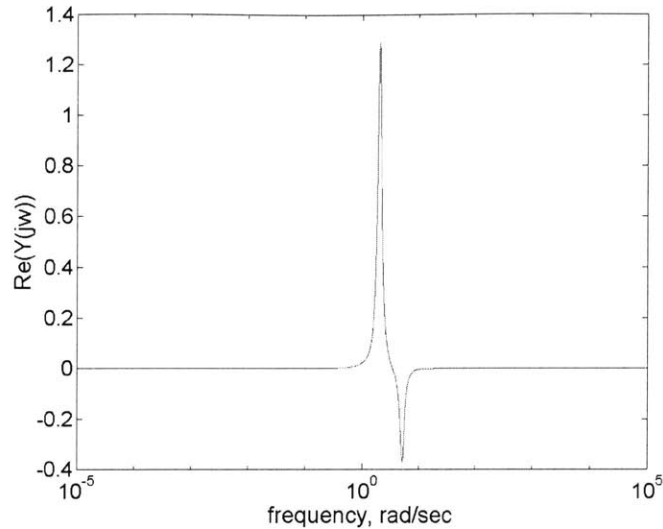
Figure 2-20: Real part of port admittance for the system in example 2.4 and 2.5.

2-20 shows the real part of the port admittance, evaluated as a function of frequency, and it is clearly negative between 2 and 10 rad/sec, hence violates the third condition for passivity.

Colgate has proven another useful result via this argument, particularly helpful in testing for coupled stability of systems. As seen in Table 1, an ideal spring in admittance causality produces $+90°$ of phase, and an ideal mass produces $-90°$, both for all frequencies, making each passive but not strictly passive. If the manipulator is non-passive, its phase must necessarily exceed either $-90°$ or $+90°$ at some frequency, so that coupling it to a pure spring or a pure mass results in phase that exceeds $-180°$ or $+180°$ at that frequency. The value of the environmental stiffness or mass acts as a gain, and can be selected such that the magnitude exceeds 1 at a frequency where the phase crosses the boundary, proving instability by the Bode criterion. Alternatively by the Nyquist criterion, the gain expands or contracts the contour, and can be selected so as to produce an encirclement of the -1 point, which must be possible if the phase exceeds $-180°$ or $+180°$. Conversely, it is impossible for the manipulator to be unstable when in contact with any passive environment if it does not become unstable when coupled to any possible spring or mass. Thus the passivity

85
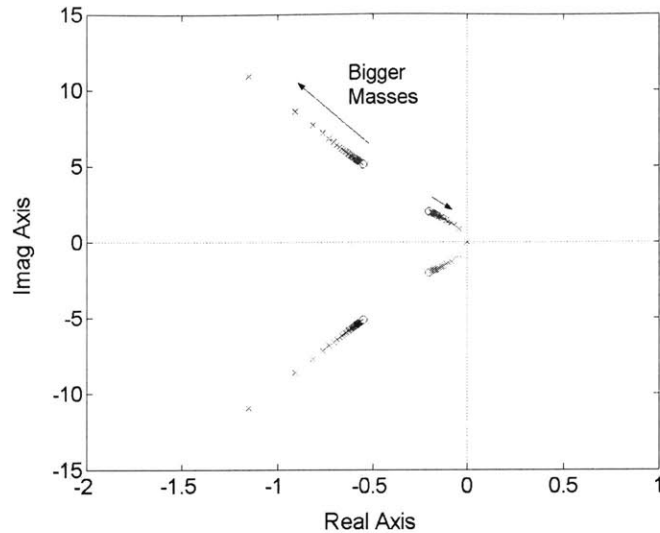
Figure 2-21: Root locus with respect to environment mass of the system in example 2.6.

of a manipulator can theoretically be evaluated by testing stability when coupled to all possible springs and masses. If no spring or mass destabilizes the manipulator, it is passive. Much can be learned about a system by understanding which environments destabilize it [23, 27].

Example 2.6: Destabilizing environments The destabilizing effect of springs or masses can be observed by examining the locus of a coupled system's poles as an environmental spring or mass is changed. Returning to the single-resonance example with force feedback, figure 2-21 shows the coupled system poles as a mass connected to the interaction port is varied from 0.1 kg to 10 kg. The system remains stable. Figure 2-22 shows the coupled system poles as a stiffness connected to the interaction port is varied from 1 N/m to 150 N/m. When the stiffness is large, the system is destabilized.
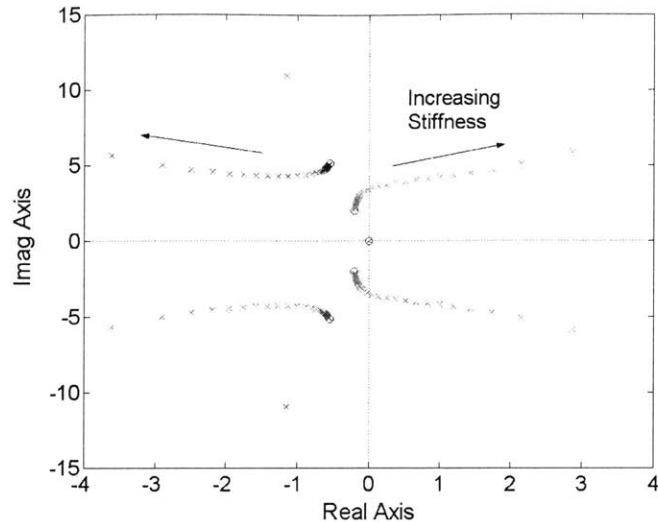
Figure 2-22: Root locus with respect to environment stiffness of the system in example 2.6.

2.3.5 Implementing interaction control

Like any control system, a successful interactive system must satisfy the twin goals of stability and performance. The preceding sections have shown that stability analysis must necessarily include a consideration of the environments with which the system will interact. It has further been shown that stability can in principle be assured by sculpting the port behavior of the system, at least when the environment is passive. Performance for systems that interact with humans (and, often, other interactive machines) is also measured by dynamic port behavior, so both objectives can be satisfied in tandem by a controller that minimizes error in the implementation of some target interactive behavior. For this approach to work for manipulation, any controller must also include a way to provide a motion trajectory. The objectives of dynamic behavior and motion can, to some degree, be considered separately.

Virtual trajectory and nodic impedance

Motion is typically accomplished using a *virtual trajectory*, a reference trajectory that specifies the desired movement of the manipulator. A virtual trajectory is much

like a motion controller's nominal trajectory, except there is no assumption that the controlled machine's dynamics are fast in comparison to the motion, an assumption that is usually required to ensure good tracking performance. The virtual trajectory specifies which way the manipulator will "push" or "pull" to go, but actual motion depends on its impedance properties and those of the environment. It is called "virtual" because it does not have to be a physically realizable trajectory.

Deviation in the manipulator's trajectory from its virtual trajectory produces a force response that depends on the interactive properties of the manipulator, its *nodic impedance/admittance*. If the robot is represented as a mass subjected to actuator and environmental forces (similar to the simple model of example 2.3 above) and the nodic impedance of the controller is a parallel spring and damper, the resulting behavior is as pictured in figure 2-23 in which v_o represents the virtual trajectory and v the actual manipulator motion. In many applications, the controller will be required to specify manipulator position x and, to be consistent, the virtual trajectory must also be specified as a position x_o. In that case, the non-inertial (nodic) behavior is strictly described as a dynamic operator that produces output force in response to input displacement (the difference between virtual and actual positions, $\Delta x = x_o - x$), which might be termed a "dynamic stiffness". However, the term "impedance" is loosely applied to describe either velocity or displacement input. To the interaction port, the system behaves as a mass with a spring and damper, connecting the port to some potentially moving virtual trajectory. The dynamics of the spring, mass, and dashpot are as much a part of the prescribed behavior as is the virtual trajectory, further distinguishing this strategy from motion control.

"Simple" impedance control

One primitive approach to implementing impedance control, proposed in [57], has been used with considerable success. Termed "simple" impedance control, it consists of driving an intrinsically low-friction mechanism with force- or torque-controlled actuators, and using motion feedback to increase output impedance. This approach makes no attempt to compensate for any physical impedance (mass, friction) in the
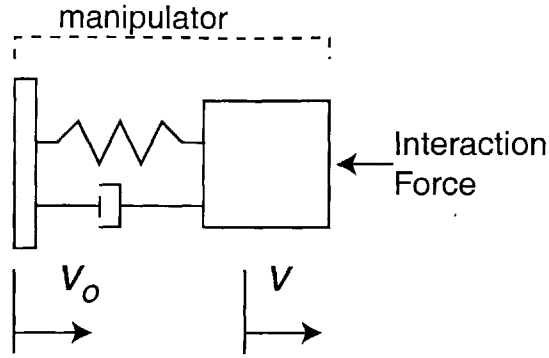
manipulator

Figure 2-23: Virtual trajectory v_o and nodic impedance of an impedance-controlled manipulator.

mechanism, so the actual output impedance consists of that due to the controller plus that due to the mechanism.

If a robot is modeled as a multi-degree-of-freedom inertia retarded by damping and subject to actuator and environmental torques (a multi-degree-of-freedom version of the model used in example 2.3 above), the robot with the simple impedance controller is as follows:

$$I(\Theta)\ddot{\Theta} + C(\Theta, \dot{\Theta}) + D(\dot{\Theta}) = T_a + T_e \qquad (2.14)$$

where Θ is a vector of robot joint variables (here assumed to be angles, though that is not essential), I is the robot inertia matrix (which depends on the robot's pose), C denotes nonlinear inertial coupling torques (due to Coriolis and/or centrifugal accelerations), D is a vector of dissipative velocity-dependent torques (e.g., due to friction), T_a is a vector of actuator torques and T_e a vector of environmental torques. The target behavior is impedance; if the behavior of a spring with stiffness matrix K_j and a damper with damping matrix B_j is chosen, the control law is simply:

$$T_a(\Theta, \dot{\Theta}) = K_j(\Theta_o - \Theta) + B_j(\dot{\Theta}_o - \dot{\Theta}) \qquad (2.15)$$

where Θ_o is a virtual trajectory in robot joint space. Combining the controller impedance (equation 2.15) with the robot dynamics (equation 2.14), the result is

as follows:

$$I(\Theta)\ddot{\Theta} + C(\Theta, \dot{\Theta}) + D(\dot{\Theta}) + B_j\dot{\Theta} + K_j\Theta = B_j\dot{\Theta}_o + K_j\Theta_o + T_e \qquad (2.16)$$

This controller implements in robot joint space a dynamic behavior analogous to that depicted in figure 2-23; the controller spring and damper serve to push or pull the robot pose towards that specified by the virtual trajectory.

Due to the nonlinear kinematics of a typical robot, the controller described above produces position-dependent endpoint stiffness and damping when constant K_j and B_j matrices are selected. This can be remedied using simple kinematic relationships that define the robot's behavior, and a nonlinear control law can be derived that uses transformations that are well-defined, even near or at the robot's singularities. Details are available in [57].

Note that the simple impedance controller does not rely on force feedback and does not require a force sensor. If the forces due to the robot's inertia and friction are sufficiently small, the robot's interaction port behavior will approach the desired impedance. Inertial forces decline for slower movements, and vanish if the robot is at a fixed pose. As a result, a simple impedance controller can be quite effective in some applications. Frictional forces may also decline for slower movements but, especially due to dry friction, need not vanish when the robot stops moving. This is one important reason why, for applications requiring low impedance, low inertia and friction are desirable in the design of interactive robots.

From a controller design viewpoint, the simple impedance control law of equation 2.15 closely resembles a PD motion controller acting on the error between actual and virtual trajectory in joint space. Given the assumptions outlined above (a robot modeled as a mass driven by controllable forces) the impedance controller stiffness K_j and damping B_j correspond respectively to proportional and derivative gains.

The simple impedance controller has another important feature. Insofar as the controller exactly mimics the behavior of a spring and damper, the robot behaves exactly as it would with a real spring and damper connecting it to the virtual trajec-

tory. If the virtual trajectory specifies a constant pose, the entire system (robot plus controller) behaves as a collection of masses, springs, and dampers, and as such is passive, guaranteeing stable interaction with all passive environments. Regardless of the actual robot mass, damping, and controller gains, coupled instability is completely eliminated. In fact, this is true even if the interaction occurs at points other than the end-effector, whether inadvertently or deliberately (e.g., consider using an elbow or other body part to manipulate) [59]. In addition, because the simple impedance controller does not rely on force feedback control to shape impedance, it is not vulnerable to the loss of passivity that can occur when structural modes interact with a force feedback loop (as illustrated in example 2.4).

Thus, although it is primitive, a simple impedance controller goes a long way towards solving the stability problem and its performance gets better as the inherent robot impedance is reduced. In practice, though this implementation performs well in some applications, it has limitations. Several factors make the controller impedance non-passive, including discrete-time controller implementation and unmodeled dynamics between actuator and sensor. Under these conditions stability cannot be guaranteed with all environments. At the same time, the creation of low-impedance hardware can be difficult, particularly for complex geometries. Still, the approach has been quite successful, particularly when used in conjunction with highly back-drivable designs (for example with MIT-MANUS, described in section 1.6). Feedback methods to reduce the appearance of intrinsic robot impedance are discussed below.

2.3.6 Improving low-impedance performance

A large class of applications, including robots that interact with humans, demands interactive robots with low mechanical impedance. The most direct approach is to design low-impedance hardware and use a simple impedance control algorithm; in fact, this is the recommended approach. However, intrinsically low-impedance hardware can be difficult to create, particularly with complex geometries and large force or power outputs. Most robotic devices have intrinsically high friction and/or inertia and the simple impedance control technique described above does nothing to compensate

for intrinsic robot impedance. Discussion in the first part of this chapter explains why it is difficult to create low-impedance robots for complex configurations. Several types of actuator technologies, however, including geared electromagnetic and servo-hydraulic, satisfy force density requirements but are ineffective because of their high intrinsic impedance (see figure 2-10). Should their apparent impedance be artificially reduced, such technologies would be viable for interaction control. Considerable effort has been devoted to designing controllers to reduce the apparent endpoint impedance of interactive robots.

Force feedback and force control

Force feedback is probably the most appealing approach to reducing apparent impedance. If a simple impedance controller is applied to a 1-D inertial mass as shown in figure 2-24, subject to some nonlinear friction force $F_f(x, \dot{x})$, and is augmented with a proportional force feedback controller, the control law is given by equation 2.13, repeated here for convenience.

$$F_a = K(r - x_1) + B(\dot{r} - \dot{x}_1) + K_f(F_e + K(r - x_1) + B(\dot{r} - \dot{x}_1)) \qquad (2.17)$$

K and B are the scalar desired stiffness and damping, respectively, and r now represents the virtual position (The motion control portion of the law is equivalent to a PD controller with position gain $K(1 + K_f)$ and velocity gain $B(1 + K_f)$. This form was chosen in order to preserve the value of the stiffness K and damping B). The force feedback term serves to minimize the deviation of the actual endpoint force from the desired endpoint force, which looks like a damped spring characterized by K and B connected to the virtual trajectory. The equation of motion for the uncontrolled system with nonlinear friction is:

$$m\ddot{x} + F_f(x, \dot{x}) = F_a + F_e \qquad (2.18)$$

92
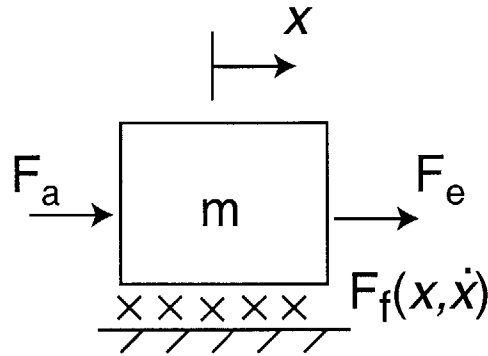
Figure 2-24: Single-mass, single degree-of-freedom robot model with nonlinear friction.

Combining equations 2.17 and 2.18, the controlled equation of motion is as follows:

$$\frac{m}{1 + K_f}\ddot{x} + \frac{F_f(x, \dot{x})}{1 + K_f} + K(x - r) + B(\dot{x} - \dot{r}) = F_e \qquad (2.19)$$

Equation 2.19 shows that the introduction of proportional force feedback reduces the apparent mass and friction by a factor of the force feedback gain plus one $(1 + K_f)$, so that (in principle) arbitrarily large force gains drive the actual impedance to the desired impedance specified by the controller. Friction, inertia or any other behavior, whether linear or nonlinear, is minimized.

Example 2.7: Force feedback control on a screw-driven robot The screw-driven robot of example 2.2 is backdrivable but has unacceptable levels of friction and inertia. One model of this system is given by equation 2.18, where the nonlinear friction can be approximated by a model $F_f(x, \dot{x})$ (a constant gravity force should also be included, but this can be directly cancelled by a constant control term and is omitted for simplicity). A simple impedance controller to implement spring stiffness K and linear damping B on this system, with the virtual trajectory fixed at $r = \dot{r} = 0$, takes the form:

$$F_a = -Kx - B\dot{x} \qquad (2.20)$$

When this control is applied to equation 2.18, the resulting controlled equation of motion is:

$$m\ddot{x} + F_f(x, \dot{x}) + B\dot{x} + Kx = F_e \tag{2.21}$$

The unwanted mass and friction appear in the controlled system dynamics. To reduce their influence in the port impedance, the force feedback law of equation 2.17 can be applied. With the virtual trajectory fixed as above, the resulting equation of motion is:

$$\frac{m}{1 + K_f}\ddot{x} + \frac{F_f(x, \dot{x})}{1 + K_f} + Kx + B\dot{x} = F_e \tag{2.22}$$

Both controllers were implemented on the screw-driven vertical robot, with target stiffness $K = 200$ N/m, zero target damping $B = 0$, and $K_f = 5$. Figure 2-25 shows the force versus position of the screw module with both controllers applied, as well as the desired behavior, a stationary 200 N/m spring. The plots were generated by slowly applying force by hand at the endpoint. Because of static friction, a large breakaway force is required to initiate or reverse movement with simple impedance control (equation 2.20). With the force feedback controller (equation 2.17), friction is reduced as expected, by roughly a factor of $K_f + 1 = 6$. Additional tests show that apparent inertia is also reduced as expected [15].

This one-degree-of-freedom combination of force and motion feedback is readily extended to multiple degrees of freedom. Assuming the robot may be modeled as a multi-degree-of-freedom inertial mechanism, retarded by friction and subject to actuator and environmental forces (as in equation 2.14 above) a controller based on feedback of endpoint position, velocity, and force can be formulated to replace the manipulator's inherent dynamics with arbitrary inertia, stiffness, and damping [57].

Despite its appeal, force feedback control has fundamental limitations. Clearly, the largest possible force feedback gain is desirable to minimize unwanted components of
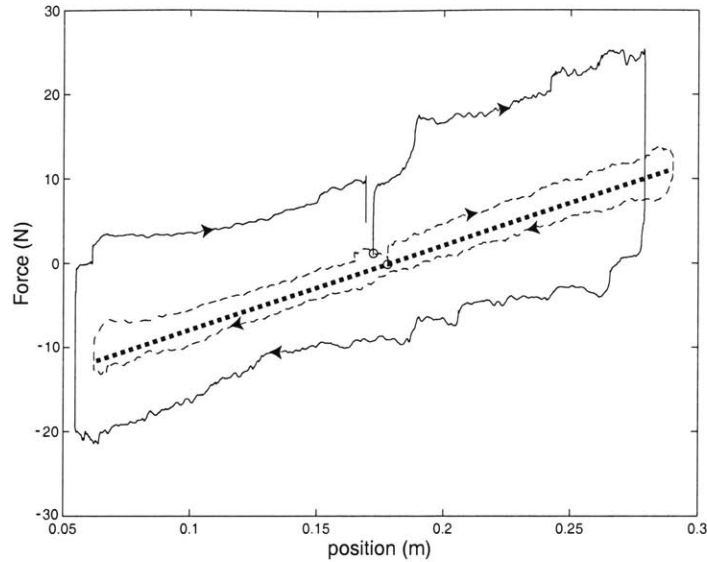
94

Figure 2-25: Virtual 200 N/m stiffness represented on the screw-driven robot from example 2.2 with simple impedance control (solid line) and proportional force feedback control with $K_f = 5$ (light dashed line). Actuated by hand. Ideal behavior shown by heavy dashed line.

intrinsic robot behavior, for example due to nonlinear friction. However, experienced control system designers intuit that increasing gains toward infinity almost always leads to instability, and the following observation shows why this holds true in the case of force feedback. Colgate has derived a *physical equivalent representation* of force feedback. A physical equivalent representation is a model of a mechanical system that exactly replicates the behavior of a system under feedback control. Figure 2-26 shows a physical equivalent representation for a model of a robot with a single structural vibration mode (as shown in figure 2-12 but with $b_1 = b_2 = 0$, $m_1 = m_2 = m/2$) under force feedback control (equation 2.17 with $K = B = 0$). It also exhibits a single structural mode but with parameters that depend on the force feedback gain. The result has negative stiffness, mass and damping parameters for any force feedback gain greater than 1 [23]. Because endpoint force feedback is not state feedback, isolated stability is unaffected; this system remains neutrally stable as long as the interaction force F_e is zero (indeed in this case, the actuator command is zero). However, for $K_f > 1$ it is not passive; it can exhibit negative visco-elastic behavior upon contact,

95
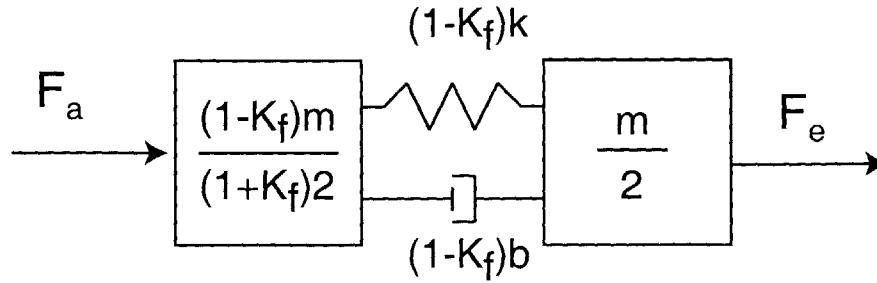
Figure 2-26: Physical equivalent representation of single-resonance robot under proportional force feedback control with gain K_f. Reproduced from [23].

resulting in coupled instability. Note that this gain limit is independent of the value of the structural stiffness and damping. As any real robot almost inevitably exhibits resonance, this analysis shows that any system under force feedback becomes non-passive for a force feedback loop gain greater than unity and is therefore vulnerable to coupled instability.

As impedance and force regulation are closely related, the extensive force control literature can be used to help understand stability properties. When the target impedance is zero and a reference force trajectory is independently specified, the force control problem results. Practically, the issues that threaten stability for force control and for impedance control methods that use force feedback are very similar; in fact several works have shown the equivalence of the two approaches (see for example [32]).

In one of the earliest widely-cited papers on force control, Whitney [118] discusses continuous force feedback as a way to assist manipulators in assembly of mating parts. Specifically, a controller that provides velocity command proportional to measured force is analyzed. The environment is modeled as a spring, and the block-diagram analysis implies the assumption that the spring does not affect manipulator motion (i.e. the environment is soft). A model with no structural resonance is used. Whitney models a sampled-data implementation of the controller and derives a stability condition: the product of the environment stiffness and force feedback gain must be less than some constant, defined by the sampling rate and robot actuation bandwidth.

The force feedback gain, and therefore performance, are limited by the environment's compliance.

In a subsequent paper published a decade later, Whitney [119] reviews progress in the field and analyzes discrete implementations of stiffness and damping control with force feedback. For stiffness control with a stiff environment, the robot must have high stiffness to achieve high force feedback gains. But if force control is the objective, a low stiffness is desired. The result for damping control from [118] is repeated, with the observation that low damping gain produces sluggish behavior and low velocities. Whitney suggests adding passive compliance to the robot end-effector to make the environment appear compliant to the robot. This strategy is thoroughly explored in chapter 5.

Eppinger and Seering [38] analyze several potential sources of instability in force feedback systems. They model the environment either as a spring-mass-damper or a rigid constraint, in which case the force transducer is modeled as a damped spring that dominates the environment impedance. Modeling the robot as a rigid body, limitations in actuator bandwidth are shown as a possible threat to instability when force is fed back. Dynamics in the robot base or the workpiece can also compromise stability if there is insufficient phase margin, even though the actuator and force sensor are collocated. If a structural resonance is included in the robot model, leading to non-collocated actuation and force sensing, stability of the coupled system is threatened at quite low gains. Of the conditions analyzed, this is by far the most significant stability risk. Although no prediction of the limiting stable gain is provided, this analysis verifies Colgate's important result: resonant dynamics between the sensor and actuator lead to potential lost coupled stability at low force feedback gains.

Despite the importance of this conclusion, it has not fully permeated the force control literature. Many recent works treat coupled stability only in passing, using collocated robot models without structural resonance to design controllers. For example, see [22, 107, 84, 115]. Many current authors report only environments with relatively low stiffness to demonstrate their controllers, avoiding the contact problems most strongly associated with stiff environments and focusing on force regulation and

tracking performance. Contact problems due to impact with stiff surfaces are also sometimes avoided by ensuring that the manipulator velocity is negligibly small as it transitions from free movement to contact [22]. This approach is common particularly in schemes that switch between different controllers for free movement and interaction, and that often require an accurate model of the environment to ensure stability and/or performance.

The addition of virtual damping, either via PD feedback of force or via velocity feedback, can help to slow velocity at contact and is also commonly reported to help stabilize coupled robot-environment systems under force control [22, 96, 77]. In [96] it is reported that virtual damping from force differentiation can permit high feedback gains in contact tasks with stiff environments.

Wen and Murphy [116] present an analytical model of force control stability on a general robot arm model. Although the model they use has collocated sensors and actuators, they find several sources of instability. One is due to the time delay in force measurement, when the environment is infinitely rigid. Although any time delay in the control loop is indeed likely to cause problems, their analysis depends on the fact that the delay introduces an algebraic loop between the measured force and the applied force, a consequence of the assumption of collocation. They propose integral action as a remedy for this stability problem, but note that high integral gains can become problematic if the environment has a small degree of compliance. It is likely that both time delay and integral control action can compromise passivity and cause coupled instability, along with structural resonance

Alternative impedance control implementations

The force feedback law given in equation 2.17 is not the only way to use endpoint force information to implement impedance control. An alternative method also uses a model of the robot to compute the desired control [57, 60]. This eases the implementation of target inertia, in addition to target stiffness and damping, without the measurement or computation of endpoint or joint acceleration. Similarly, the advantages of a model-based approach in implementing desired dynamic behavior (as

opposed to desired steady-state behavior) have been observed in [107].

A different approach to implementing impedance control is to use a fast, inner motion control loop in conjunction with an outer force control loop [77, 107, 32]. This approach can in some cases leverage some of the advantages of high-bandwidth control, which may be achievable in the motion loop but not the force loop because the motion loop may be collocated. In this architecture, the force loop is usually used to determine the desired motion to provide port behavior that approaches that of the target impedance in response to the applied force history. The dynamics of this loop are fixed by the impedance specification. The output of this loop is a motion command, which serves as the input to the motion loop. The motion loop then servos to this trajectory by a preferred method. Certain disturbances can in theory be rejected more effectively by this fast servo loop. Of course, the interaction problem remains, and motion tracking stability is insufficient to guarantee coupled stability. An example of a method that uses an interior motion loop to achieve passive and stable interaction is provided in the next section.

Strategies that use an interior motion servo loop are especially appealing for robot systems that are designed for motion regulation, including almost all industrial robots. This permits the use of the motion-control technology for interaction control. When low-impedance hardware is designed, simple impedance control is preferred as it avoids the destabilizing effects of force feedback. When custom hardware with moderate impedance is designed, the preferred method is less obvious.

Natural admittance control

One effective solution to the loss of passivity due to force feedback is natural admittance control (NAC) developed by Newman [87]. In essence, this approach is based on the observation that a system under force feedback becomes non-passive when the controller reduces the apparent inertia (to less than half its physical value) but passivity is not compromised by the elimination of friction. Natural admittance control specifies a desired inertia that is close to that of the physical system (the "natural admittance") and focuses on reducing friction as much as possible, preserving

99

passivity.

A method for the design of natural admittance controllers is detailed in [87]. Here the procedure is sketched for the simple robot model shown in figure 2-11 consisting of a single mass retarded by friction and driven by actuator and environmental forces. The system might also have nonlinear friction that the NAC seeks to eliminate, that need not be modeled at all in the design of the compensator; the controller treats such friction as a disturbance from desired port behavior, and rejects it. To ease notation, the substitutions $v = v_a = \dot{x} = v_e$ are made. The equation of motion for the system velocity neglecting nonlinear friction is:

$$mv + bv = F_a + F_e \qquad (2.23)$$

In the Laplace domain:

$$v = \frac{1}{ms + b}(F_a + F_e) \qquad (2.24)$$

A generic form for the controller is assumed, that incorporates some velocity feedback with compensator $G_v(s)$ and endpoint force feedback with some compensator $G_f(s)$:

$$F_a = G_v v + G_f F_e \qquad (2.25)$$

Using equations 2.24 and 2.25 the actual endpoint admittance $Y(s) = \frac{v}{F_e}$ is determined:

$$Y(s) = \frac{G_f + 1}{ms + b - G_v} \qquad (2.26)$$

This is equated to the target port admittance. The target stiffness K and damping B are chosen at will but the target mass is equal to the physical mass m of the system:

$$Y(s) = Y_{des}(s) = \frac{1}{ms + B + \frac{K}{s}} \qquad (2.27)$$

Equations 2.26 and 2.27 can be solved for G_f in terms of G_v:

$$G_f = \frac{(b - G_v - B)s - K}{ms^2 + Bs + K} \qquad (2.28)$$

100

A simple form for G_v can be assumed, such as a constant, and the compensator design is complete. Equation 2.28 may be thought of as a force feedback "filter" that yields the desired admittance of equation 2.27. Although the mass is not reduced, the "disturbance" to the desired behavior due to friction, b, is rejected with the feedback gain of the compensator, and its effect is smaller as the velocity loop gain, G_v, is increased. The approach serves equally well to minimize disturbances due to nonlinear frictional forces, e.g., dry friction.

A similar formulation can be used to increase or decrease the target apparent inertia. In theory, increasing the inertia should not pose a stability risk. Decreasing the apparent inertia to less than 50% of the actual value necessarily makes the system nonpassive.

In principle, the controlled system is passive even if the velocity gain is increased to arbitrary size, minimizing the unwanted frictional effects. In practice unmodeled dynamic effects limit the gain that may be used without compromising passivity, but the technique affords significant practical improvement over simple proportional force feedback. A more general formulation and discussion of the method can be found in [87, 88].

Time-domain control strategies

Several alternative methods, based in the time domain rather than the frequency domain, have been proposed for stable interaction control. Gillespie [47] uses a second-order model of the human limb (a finger) with a model of a haptic device and a virtual wall to predict time histories and compensate for "energy leaks" produced by nonideal control behavior. A controller adds to the base haptic controller at each sample when necessary to compensate for destabilizing behavior resulting from discrete sampling. A "watchdog function" is proposed to track problems that arise from switching the wall on and off as the finger moves into and out of contact with it. A deadbeat controller is used to compensate for these effects as they arise. This technique requires a parametrized model of the environment, and specifically targets two problems that result from sampled-data implementation of the virtual wall.

Hannaford [52] takes a passivity-based time-domain approach. He proposes monitoring port variables with a "passivity observer" to measure whether the total energy output exceeds the total energy input over time. When active behavior is found, he proposes a "passivity controller" that drains the excess energy that makes the system appear active. Several different methods are suggested to implement the passivity controller as part of the virtual coupling in the haptic system. Simulation and experimental results are presented that show this approach stabilizing a very high interface stiffness and compensating for control system delay. This approach requires no model and is computationally light.

The appeal of these approaches is that they theoretically operate only when needed, when coupled stability (or passivity) is jeopardized. They do not affect performance when stability is not at risk.

Stramigioli has proposed "intrinsically passive control," a hierarchical approach that guarantees system passivity by guaranteeing passivity of individual system parts, which are interconnected by power bonds [109]. This permits the separation of the low-level and high-level portions of the controller without sacrificing the passivity of the controlled system. The example of the low-level controller provided in [109], however, simply uses proportional-derivative motion feedback, and does not help reduce apparent robot impedance. Presumably it could be replaced with a more advanced passive control strategy.

Passive control limitations

Thanks to its appeal as an intuitive, energy-based method, passivity remains the standard approach for interaction stability. However, the preceding sections have highlighted several limitations and difficulties associated with controlling systems to be passive.

Passivity is a difficult standard to meet in practice; a virtually limitless number of unmodeled phenomena are available to sabotage the designer's best intent, from discrete control, sensor limitations and time delays to nonlinear friction, resonant modes, and backlash. Even after control is implemented, it is challenging to be certain

that a system is passive. Colgate shows that verifying stability with ideal springs and inertias of all possible magnitudes is in principle sufficient to show passivity [23], but in practice, this procedure could at best be approximated.

Even an ideally implemented passive controller ensures stability only with passive environments, and offers no assurance if the environment is the least bit nonpassive. However, the human neuromuscular system uses sensory feedback loops and includes substantial delays, characteristics which often render physical systems nonpassive. Historically, passive robots have successfully interacted with humans without incident, and this has been justified on the basis of limited energy production or on the limited frequency range of active human motion production [56]. Strictly speaking, however, this is not a perfect fit to the theory.

Finally, passive controllers may be needlessly limiting in certain important respects. As noted above, a force feedback controller renders a robot nonpassive when it reduces the apparent inertia below half of the actual value. This guarantees that the machine couples unstably to some passive environment. But this statement says nothing about the characteristics of the destabilizing environment. Perhaps the set of destabilizing environments for a particular robot and controller is limited to dynamics that are not present in the robot's sphere of interaction, and though the controller makes the robot nonpassive, this may be satisfactory. Yet if passivity is used as the stability measure, the controller must be made less aggressive (from which performance may suffer). Significantly, passivity is strictly a phase constraint on the port function; the magnitude of the admittance or impedance of any passive system can be arbitrarily small or large. Thus if stable interaction with a small mass is achieved through passivity, stability is also achieved when interacting with a stiff spring. If the robot will encounter no spring-like environments, passivity may not be a suitable measure of stability.

Passive control design is completely independent of the environment. This is probably its greatest advantage, but also perhaps its greatest limitation. This can produce controllers that are insufficiently conservative (if the environment is nonpassive), or that are excessively conservative (if the environment has limited dynamic

properties). In the following chapters, an alternative measure of coupled stability is proposed that incorporates limited knowledge of the environment to shape controllers that take advantage of the properties of both interacting systems, rather than just one.

Chapter 3

Design of Controllers Using Limited Knowledge of the Environment

Designing a stable interaction controller amounts to synthesizing a controller that stabilizes the coupled system, consisting of the controlled robot and the environment. As defined in chapter 1, the stability condition must hold for all environments the robot is expected to encounter. Subject to this constraint, control should be adjusted to improve performance. While considerable study has been devoted to manipulating robot port behavior, to date little has been done to include the characteristics of the environment in the design of controllers for interactive machines. Here the following hypothesis is explored: the better defined the constraint that coupled stability imposes, the less conservative controllers can be made, and the better performance may become. This and the following chapter consider this hypothesis, and trace the development of a new method to design controllers for interactive machines that improve performance by exploring a new region of the design space, using limited knowledge of the environment. Machines with these controllers are said to be *complementary* to the environments with which they are designed to interact; that is, they stably interact with these particular environments. A passive robot is complementary to the set of passive environments. A robot need only be complementary to the set of envi-

ronments with which it will interact; coupled stability with unexpected environments is irrelevant to the definition of stability provided in chapter 1.

An alternative stability measure is proposed for the case when limited knowledge of the environment is available. Thanks to a large body of human subject research, this is the case for human limbs. A systematic method of controller design is proposed to optimize robot performance subject to the stability constraint. In the next chapter, the method is applied to an example and improved performance is demonstrated versus passive control.

3.1 Passive control and human-robot interaction

Passive controllers represent the state of the art in interaction control. Yet as noted in chapter 2, these controllers leave a great deal to be desired. While they can, at least in theory, reject friction while providing stable interaction, they are unable to reject more than half of a system's intrinsic inertia. While this sets a hard limit on performance, passivity is also a questionable choice as a stability measure for coupling to human limbs.

3.1.1 Passivity as too conservative for human limb interaction

When passive controllers are designed, no data, neither model- or behavior-based, is used to characterize the environment. The only assumption made about the environment is that it is passive, i.e. that it will not generate energy. The set of passive environments includes systems with such diverse dynamics as a kinematic constraint, a massive workpiece moving freely, a massless spring of any stiffness, an interface dominated by static friction, or a highly nonlinear damper (as well as many more). None of the potential environments on this list behave like human limbs. Human limbs have limited, and often well understood and modeled, mechanical impedance properties. In particular, the magnitude of the mechanical impedance is limited. Even consider-

ing a wide range of subjects, human limbs have minimum and maximum mass, and maximum stiffness. No human limb is ever massless, infinitely rigid, or dominated by static friction. Studies have been done to characterize the mechanical properties of various human limbs, joints, and muscles; examples include [65, 49, 50, 92, 93, 91]

If the magnitude of human limb impedance is bounded significantly, then passivity may be exceedingly conservative for human-interactive robots, and performance may be sacrificed in achieving this objective. Because interactive robots are designed to have dynamics comparable to those of the human limbs with which they interact, it is likely that the impedance magnitude bounds are, in fact, significant. Thus to the robot, an extremely stiff human arm is likely to appear very different from a kinematic constraint. Requiring coupled stability in interaction with a kinematic constraint, as is required when passivity is the stability measure, is completely unnecessary and exceedingly conservative. In contrast, the approach proposed here incorporates the known bounds of the environment's impedance magnitude, in addition to its phase behavior, to predict coupled stability. A procedure to optimize performance among stable solutions is included.

3.1.2   Passivity as insufficient for coupled stability

Strictly speaking, human limbs are likely nonpassive, as the neuromuscular control system includes feedback and large known time delays. As noted in the previous chapter, several arguments have been made that this does not provoke instability with passive robots, and in practice this has generally proven true. Nevertheless, the approach proposed here allows active environments, provided that an accurate model or complete data are available to describe them. With this method, the robot controller can be sculpted to stabilize contact even with non-passive environments. Should this be necessary for the human interaction problem, it can be accommodated in a straightforward way.

3.2 Non-passive robots in practice

In addition to the intuition that relaxing passivity should permit improved performance via more aggressive control, there is empirical evidence for this as well. Force feedback is used in many force-controlled and impedance-controlled robots in real applications, where passivity is seldom achieved. Force feedback gains are usually tuned manually to an appropriate stability margin, based on testing. The result is not necessarily passive. An example is provided.

> **Example 3.1: Non-passive screw robot control** As detailed in example 2.7, a force feedback controller was implemented on the screw-driven vertical robot module in a bid to reduce its apparent impedance. Extensive testing with human arms revealed the coupled system to be stable at gains exceeding $K_f = 5$. The controller successfully reduced apparent Coulomb friction and inertia by a factor of 6. With proportional force feedback gain exceeding 1, this system is nonpassive. Yet it has been used effectively in clinical trials with this controller [74]. Contact with human arms never results in a coupled system instability.

If passivity is lost when the force feedback gain exceeds 1, how can the system in this example function with a gain five times larger? The answer lies in the properties of the environment. The system is, of course, nonpassive; when it contacts stiff springs, the coupled system is unstable. But this problem never arises when interacting with human arms, because human arm dynamics differ substantially from stiff massless springs. By at least one measure of performance, the apparent inertia, this nonpassive (yet stable) controller outperforms any possible passive compensator. By reducing inertia by a factor of six, performance by this measure is three times better than with passive control.

Were this system required to interact with stiff massless springs, this control would not suffice; it is the environment's dynamics that permit more aggressive control. Any control design approach that neglects to use information about the environment cannot possibly predict a controller with these properties.

3.3 Classical control design with model of environment

The background in chapter 2 establishes passivity as the standard for ensuring coupled stability. Yet the preceding example shows that a nonpassive system can couple stably to its environment and exceed performance of the same physical system under passive control. Despite the fact that it is nonpassive, this system exploits a characteristic of the environment (in this case, that it has limited stiffness) to expand the space of stable interaction controllers. As in much of the force control literature, the most problematic environments are simply avoided, permitting more aggressive force feedback. The design of the control law in equation 2.17, however, is completely ignorant of the environment. Instead, the controller is formed from a simple architecture that is known to improve performance (albeit at the expense of passivity), and coupled stability is determined strictly empirically.

If information (e.g. a model or test data) about the environment is included in the interaction controller design process, two benefits are likely: first, conditions for coupled stability and instability can be predicted rather than simply tested, improving confidence in the results, and second, performance can be improved over arbitrary designs like in the previous example. The second benefit is expected, though not assured, because additional information should provide the foundation for better control.

If the robot and environment are each represented with their port functions, one an impedance Z and the other an admittance Y, the coupled system can be modeled as the unity negative feedback system shown in figure 3-1. If the environment is exactly known and modeled by its port function, then coupled stability is equivalent to stability of the loop in figure 3-1. The passivity approach focuses entirely on constraining the phase of the robot port function alone. This neglects fully half of the coupled system. Below, a controller design approach is proposed that establishes stability on the basis of the entire coupled system, and performance on the basis of the robot port function alone.
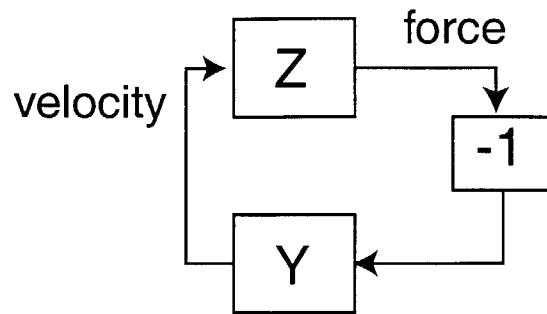
Figure 3-1: A coupled robot and environment represented by their port functions as a unity negative feedback system.

3.3.1 Servo control and interaction control

Because the design of controllers for unity negative feedback systems is so well-traveled, it is worth considering how much of this expertise can be applied to the interaction problem, specifically when a full model of the environment is available. While the *analysis* tools from classical control theory can be used effectively for the interaction problem, structural differences between servo and interaction control prevent the complete use of the *design* tools.

Stability

Stability of the loop in figure 3-1, and therefore of the coupled system, requires that the closed loop system poles be located in the Laplace domain left-half plane (with poles on the $j\omega$ axis denoting marginal stability). For a unity feedback system the location of the closed loop poles is dictated entirely by the open-loop transfer function, in this case ZY (or YZ) . The Bode and Nyquist methods provide guidelines to check the open-loop transfer function for stability of the closed-loop system [45].

This matches closely the analysis of a servo system, as shown in figure 3-2. In this system, $G_c(s)$ is the controller and $H(s)$ is the open-loop plant. Closed-loop stability is dictated by the properties of the open-loop transfer function G_cH (specifically the roots of $(1 + G_cH)$, and can be analyzed with the Bode or Nyquist criteria.

Stability margins have similar meaning for the two problems, but are more clearly

110
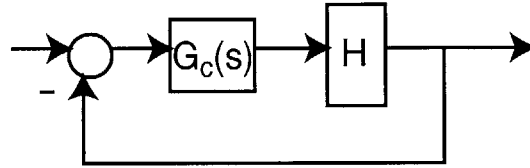
Figure 3-2: Block diagram representation of a typical servo-controlled system.

tied to controller properties for the servo system. The gain margin, for instance, is defined as the inverse of the magnitude of the open-loop transfer function at the frequency where the phase crosses $\pm\pi$ radians [45]. In the servo system, gain at this frequency can be directly increased or decreased by adjusting the loop gain in the controller $G_c(s)$. In the interacting system, however, the loop gain is not influenced directly by the controller gain, because the controller does not enter the open-loop transfer function linearly. This point is made explicitly in a subsequent section. The result, however, is apparent: while classical tools can be applied perfectly to ascertain stability and stability margin, they cannot be used as in servo design to stabilize an unstable system, or to increase or decrease stability margins. That is, they are useful for analysis, but not for design.

Performance

For a servo system, performance may be measured by a number of frequency- and/or time-domain requirements (e.g. rise time, settling time, overshoot, steady-state error, bandwidth, damping ratio, etc.). All of these are properties of the closed-loop system, the poles of which are defined by the roots of $(1 + G_cH)$. This leads to the beauty of classical servo controller design: the same quantities determine both performance and stability. Phase margin and gain margin provide measures of the stability margin as well as approximate information about performance [45]. Alternatively, a root locus can be used to locate the closed-loop poles such that they are stable and provide the desired response properties. Unlike stability, performance also depends on the

111

location of the closed-loop zeros, but these also derive directly from the open-loop transfer function. The stability and performance problems are represented in nearly identical "design spaces" and, to some degree, are solved together.

For an interactive system, particularly a human-interactive system, the story is quite different. As noted in the first chapter of this thesis, the "feel" of the robot is a key measure of performance. The impedance Z_{robot} (or its inverse) is one mathematical representation of feel. This property and any associated performance measure are completely independent of the environment Y_{env}. Performance depends exclusively on the behavior of only a portion of the open-loop transfer function. In other words, performance is a function not of the coupled, closed-loop system dynamics, but of only the robot's dynamics presented to the environment at the interaction port. These dynamics indeed play a critical role in determining stability, but they are not alone in that; the dynamics of the environment are equally critical to stability, but completely irrelevant to performance. A design method that locates the coupled system poles (such as root locus design) or that shapes the open-loop transfer function (such as Bode or Nyquist design) does not permit the appropriate performance measure to be tuned.

So, in contrast to the servo problem, any method to solve the interaction problem must either treat stability and performance as two completely separate issues, or include a new way to join the two together.

Controller-plant structure

Further comparison of figures 3-1 and 3-2 reveals an important additional difficulty in applying classical control tools directly to the interaction problem. As noted above (and in any classical control textbook), the magnitude and phase of the open-loop transfer function (G_cH) are the quantities manipulated in Bode or Nyquist control design. Through variation of the structure and parametrization of G_c, the controller, these quantities are varied simply and intuitively. The magnitude of G_cH is the
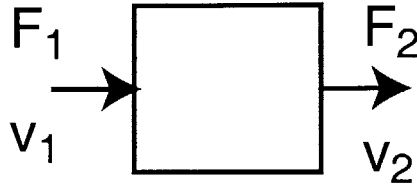
Figure 3-3: A generic two-port.

product of the magnitudes of G_c and H:

$$|(G_cH)| = |G_c||H| \tag{3.1}$$

The phase of G_cH is the sum of the phases of G_c and H:

$$\angle(G_cH) = \angle(G_c) + \angle(H) \tag{3.2}$$

This is useful because, presuming that the plant has already been designed and built, G_c is where the control designer exercises his influence. Limitations of control generally derive from the fact that the magnitude and phase of G_c cannot be varied independently at each frequency, but within those constraints the magnitude and phase of the open-loop transfer function can be altered directly by adding, removing, or changing poles, zeros, or gains of G_c. This is the case regardless of the specific method(s) the designer chooses to employ.

In contrast, no controller appears directly in figure 3-1, the interaction problem. In fact, the effect of any controller is included in the port function for the robot (e.g. Z), which also includes the model of the robot hardware. If the robot hardware is assumed to be a generic "black box" system, with an actuator driving one side and an interaction port at the other,this can be represented with a two-port as shown in figure 3-3. The force and velocity at the actuator side, F_1 and v_1 respectively, and at the environment side, F_2 and v_2, are shown in the figure. In between are arbitrary dynamics. Such a model can be used for any single-input, single-output robot.

If a linear robot model is assumed, the two-port can be defined by a matrix of

113

transfer functions relating the velocity and force at each port to each other:

$$\begin{bmatrix} v_1 \\ F_2 \end{bmatrix} = \begin{bmatrix} T_{11} & T_{12} \\ T_{21} & T_{22} \end{bmatrix} \begin{bmatrix} F_1 \\ v_2 \end{bmatrix} \tag{3.3}$$

Each of the T_{ij} entries are transfer functions. If the port variables v_1, F_2, and v_2 are fed back and each multiplied by compensators G_1, G_2, and G_3, respectively, the resulting control law is:

$$F_1 = G_1 v_1 + G_2 F_2 + G_3 v_2 \tag{3.4}$$

If this is applied to equation 3.3, the output impedance at port 2 is:

$$Z = \frac{F_2}{v_2} = \frac{T_{21} G_3 + T_{22} - T_{11} T_{22} G_1 + T_{12} T_{21} G_1}{1 - T_{21} G_2 - T_{11} G_1} \tag{3.5}$$

Note that none of the compensators G_i, and no combination of them appears linearly in the port impedance Z. Therefore no controller enters linearly into the open-loop transfer function, and none can be easily manipulated to influence the open-loop magnitude and phase. This result has significant implications for control design. If excess phase of Z is causing instability, it is not obvious how to alter the compensator(s) to remedy the problem. If stability is, in fact, achieved, the sensitivity to control parameters is far less obvious than in the case of servo design.

Multiple feedback loops

Another way that interaction control can differ from single-input, single-output (SISO) servo control is that it often includes multiple feedback loops for motion and force. Note that this is the case for example 2.7. This highlights an advantage of using port functions for interaction: the SISO loop pictured in figure 3-1 is a valid representation for interaction at any single-degree-of-freedom interaction port. This is independent of actuation and control structure. The system can use a multi-input, single-output controller, as in the example, or even a multi-input, multi-output controller, if redundant actuators are used. Regardless, the system can still be analyzed with classical

114

tools.

This observation, along with the preceding section, reinforces the fact, stated in chapter 2, that figure 3-1 provides a distinctly different way of mathematically representing dynamic systems than the usual servo structure, despite the apparent similarities. This may raise the question of whether it is ever advantageous to formulate an interacting system in the traditional, servo form. In chapter 2, the argument given against this approach is that it is invalid if coupling to the environment substantially alters the system dynamics because of loading; the implicit assumptions in using block diagram analysis are broken if the dynamics of the system are changed by interaction. In short, if a motion servo control structure is applied to a generic robot model like that in figure 3-3, the plant transfer function is $H = \frac{v_2}{F_1}$. The motion output v_2 can be validly used in a block diagram analysis only if it meets no impedance large enough to load or alter the dynamics of the transfer function. So, coupling to an environment that changes the system dynamics invalidates the use of this formulation. If the environment is completely known, and is included in the model, however, then sometimes a valid block diagram can be created. It is worth considering an example of this case to see the benefits and limitations of such an approach.

Force and motion feedback for interaction

Figure 3-4 shows a single-resonance robot consisting of two masses m_1 and m_2 joined by a spring k and driven by an actuator with force F_a, coupled to an environment that consists of a spring k_e connected on one end to the robot and on the other to ground. The position of the first and second mass are x_1 and x_2, respectively, and the interaction force between the robot and environment is F_2. If the controller feeds back x_1 and F_2, such that motion feedback is collocated and force feedback is non-collocated, the control law with generic compensators G_1 and G_2 is:

$$F_a = G_1 x_1 + G_2 F_2 \tag{3.6}$$

115
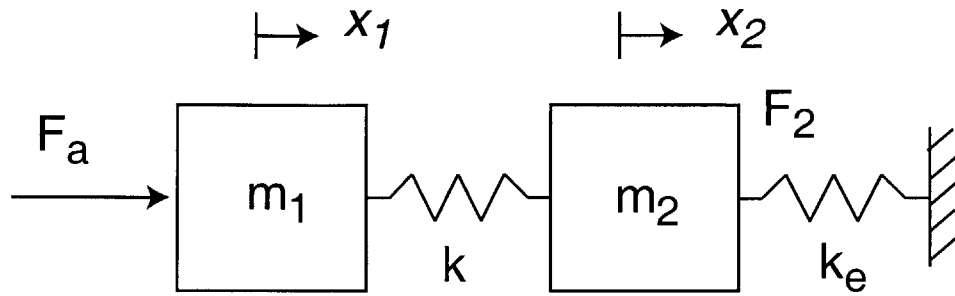
Figure 3-4: Undamped single-resonance robot model, connected to ground through the spring k_e.
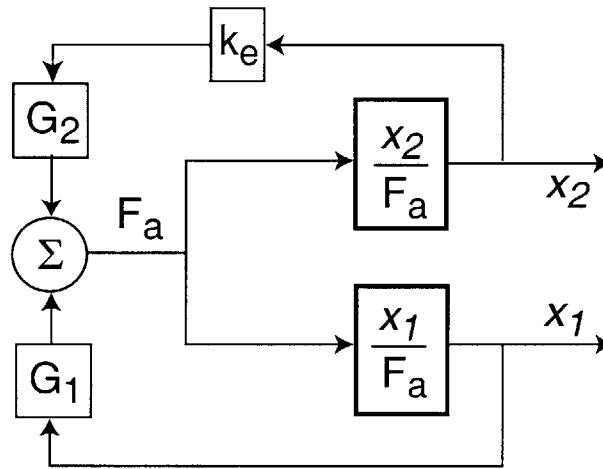
Figure 3-5: Block diagram of a single-resonance system with motion feedback (collocated with the actuator) and endpoint (noncollocated) force feedback.

Under these conditions the entire system, robot and environment, can be validly represented with the block diagram in figure 3-5, with only the assumption that the control is an ideal force source. The two transfer functions are computed from the equations of motion and equation 3.6:

$$\frac{x_2}{F_a} = \frac{k}{(M_2 s^2 + k + k_e)(M_1 s^2 + k) - k^2} \tag{3.7}$$

$$\frac{x_1}{F_a} = \frac{M_2 s^2 + k + k_e}{(M_2 s^2 + k + k_e)(M_1 s^2 + k) - k^2} \tag{3.8}$$

The operative question is: can a controller be designed for performance and stabil-

ity with this structure? There are several conditions under which such a multiple-loop system can be treated as a single-loop system and designed with single-input, single-output (SISO) tools:

- If $G_2 = 0$, the top loop is eliminated, there is no force feedback and only collocated motion feedback is used. With a single loop, stability can be ensured by applying classical SISO techniques to the design of G_1. Performance, in terms of the robot port impedance, can be addressed through simple impedance control (see chapter 2). If the environment varies but retains the same structure, uncertainty in k_e must be included. This is slightly more complicated than it is with the port function structure (shown below) because k_e appears in both the numerator and denominator of $\frac{x_1}{F_a}$, but is manageable.

- A second degenerate case arises if the motion feedback loop is eliminated with $G_1 = 0$. This amounts to a force control problem, and is in a form suggested by Eppinger to design force controllers [38]. Stability of this loop can be ensured via SISO methods. Like the preceding case, if the environment stiffness varies, uncertainty must be included in the numerator and denominator of the forward-path transfer function. As this is a force control problem, the likely performance objective is zero impedance. This can indirectly be addressed by designing the controller to have as high a magnitude as possible, to reject the internal dynamics as much as possible, while stability is maintained.

- Finally, multiple loop problems can be simplified and solved via SISO methods if all of the dynamics of one loop are sufficiently fast that they are negligible in the other loop. Here, the open-loop plants of the two loops share the same poles, so separating the dynamics of the two loops in frequency is unlikely. Appropriate feedback might render one loop much faster than the other, permitting approximate analysis, but this would require high feedback gains that are likely to threaten stability. Furthermore, such a constraint would severely limit controller design options.

In summary, although a system with an impedance controller and an explicit

117

model of the environment can sometimes be structured in a form similar to that for a servo system, this method does not allow the full use of SISO tools in the general case. In several degenerate cases the tools can be at least partially applied, but only when the impedance control problem coincides either with the motion control problem (as in simple impedance control) or with the force control problem (when target impedance is zero). Even in these cases, uncertainty in the environment is difficult to include because the environment dynamics appear mathematically in the same place as the robot dynamics. If uncertainty extends to the structure of the environment, this approach is completely invalid.

In general, SISO tools cannot be effectively applied to design interaction controllers for stability and performance together. This holds true even in the case when the environment is completely and accurately modeled, a case that never arises for human-interactive robots.

3.4 Interaction control design as optimization constrained to robust stability

As the stability and performance in interaction problems do not lend themselves to an obvious combined solution, as in the servo problem, one approach for controller design is to treat them separately, and look for a solution that satisfies both sets of requirements. All acceptable solutions must satisfy coupled stability. Amongst those that do, better solutions also improve performance.

The procedure suggested here involves generating a large number of proposed controllers. Each proposed controller must then be tested for coupled stability. A new coupled stability criterion is proposed for systems where limited knowledge of the environment is available. If a controller satisfies the coupled stability condition for all expected environments, its performance is then evaluated and compared with the performance of other controllers that satisfy the stability condition. The large set of proposed controllers is thus searched for the best solution or solutions.

The actual search (or optimization) can be performed by any number of methods. The specific steps to determine coupled stability and to evaluate performance must both be defined.

3.4.1 Evaluating stability

Limited, uncertain information is generally available about the environment for human-interactive devices. An exact model is almost always unattainable because the dynamic properties vary substantially with pose, muscle activation, and across subjects. If deviation from some nominal can be bounded, however, the environment can be modeled with a port function with some limited uncertainty. If the (scalar or matrix) environment admittance Y is modeled as a nominal admittance (scalar or matrix) function Y_n with some bounded uncertainty Δ (a scalar or matrix quantity) represented as an additive perturbation, then the following models the set of possible environments:

$$Y = Y_n + \Delta \tag{3.9}$$

The additive perturbation structure is used as an example, but this is not a requirement; multiplicative or feedback perturbations can also be used. The perturbation structure should be selected on the basis of what is known about the environment. Δ can range in value within some specific set representing the range of the expected environments' deviation from Y. The set must be constructed such that equation 3.9 can represent all expected environments using some Δ from the set. Because the uncertainty set is *bounded*, all values of Δ must obey some constraint on their magnitude. The perturbation should be constructed to satisfy the following:

$$||\Delta||_\infty \leq ||\Delta_{max}||_\infty \tag{3.10}$$

The notation in this expression indicates the infinity norm [30, 104]. Here, Δ_{max} is introduced as the instance of Δ (from within the set) with maximum infinity norm. With this additive perturbation, the coupled robot and environment is depicted

119
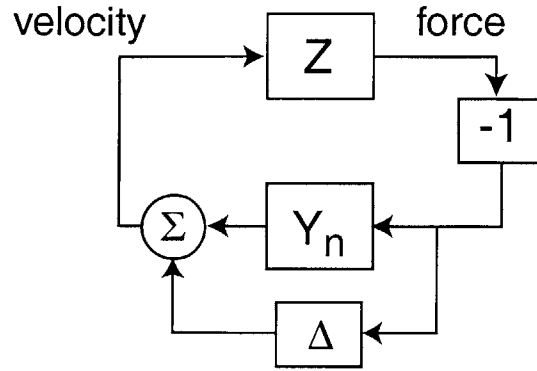
Figure 3-6: Robot interaction with a bounded set of environments.

schematically in figure 3-6, where the robot port impedance is represented by Z, and Δ.

If figure 3-6 is an accurate model of the coupled system, then *complementary stability*, a new measure of coupled stability for interaction, can be defined.

> **Definition:** A robot with port impedance function Z satisfies the *complementary stability condition* with a set of environments defined by equation 3.9 if the model in figure 3-6 is robustly stable for all values of Δ satisfying equation 3.10.

In short, if the expected environments are adequately represented by the set of admittance functions Y, then robust stability of the loop in figure 3-6 guarantees coupled stability of the robot and all expected environemnts. Analogous definitions can be introduced for alternative perturbation structures; all that is required is to replace equation 3.9 with the appropriate perturbation structure, and to adjust figure 3-6 accordingly. Definitions of alternative perturbation structures are available in [30, 104].

Nominal stability is evaluated for the case where $\Delta = 0$. If the system is nominally stable, robust stability can be determined by the *small gain theorem*.

120
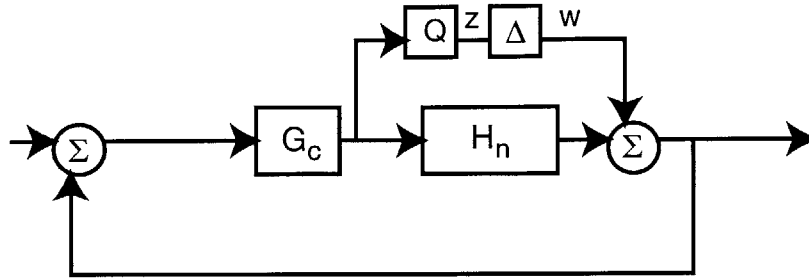
Figure 3-7: Feedback control system with uncertainty.

Robust stability and small gain theorem

If the plant $H(j\omega)$ of a feedback-controlled system contains some uncertainty such that it can be represented as a nominal plus some bounded deviation:

$$H(j\omega) = H_n(j\omega) + Q(j\omega)\Delta(j\omega) \tag{3.11}$$

where H_n is the nominal plant, Q is some shaping function, and Δ is a complex-valued function of frequency with magnitude less than one at all frequencies, then the system with feedback controller G_c is shown in figure 3-7. (Here, the additive perturbation structure is used for demonstration; other perturbation structures such as multiplicative or feedback are equally functional and do not alter the resulting theorem).

If the system is mathematically transformed to define a transfer function T_{zw} between the output from Δ, w, and the input to Δ, z, an equivalent system is pictured in Figure 3-8. Stability of this loop is characterized by the properties of the open-loop transfer function $T_{zw}\Delta$. In fact, the closed-loop characteristic polynomial for this system is $(I - T_{zw}\Delta)$. If $(I - T_{zw}\Delta)$ is nonsingular, it has no zeros in the right-half plane and, assuming that both T_{zw} and Δ are stable, the closed-loop system is stable because it has no poles in the right-half plane.

Definition: The *small gain theorem* states that the quantity $(I - T_{zw}\Delta)$ is nonsingular, and the system in figure 3-8 is therefore stable for all
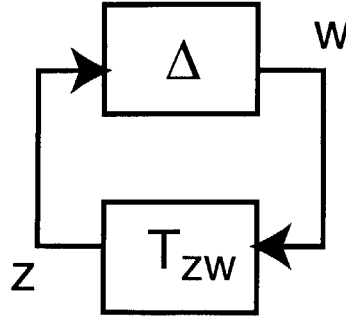
Figure 3-8: Transformed equivalent of figure 3-7, for determining robust stability.

admissible Δ , if and only if

$$||T_{zw}||_\infty ||\Delta||_\infty < 1 \qquad (3.12)$$

In general, Q is specified such that $||\Delta||_\infty \leq 1$, in which case the stability requirement is that

$$||T_{zw}||_\infty < 1 \qquad (3.13)$$

Application of the small gain theorem requires that T_{zw} and Δ be stable [30, 104].

This theorem can be applied to evaluate the stability of any system with an uncertain quantity that can be validly represented by a block diagram. Thus it can be applied to the interaction problem via structure like that in figure 3-6.

One advantage of this approach is that, depending on the actual solution technique used, the environment port function and its uncertainty can be represented either by a model with variable parameters, or directly by a data set. If a sufficiently rich set of test data is available that characterizes the full range of environment properties, the nominal and deviation from it can be numerically derived from the data, without need for an explicit model. All that is needed is a bounded region of complex numbers at each frequency.

If the environment is defined at each frequency as the set of all complex numbers with positive real part, then the set of all passive environments accurately character-

122

izes the environment. In any other case, passivity is either a conservative or inaccurate assumption. In these cases, assuming that an accurate environment model (or data set) is available, the proposed measure of complementary stability is more accurate and therefore a more appropriate stability measure than passivity.

Prior application of robust stability to coupled stability

Several earlier works have employed robust stability tools in designing controllers for interaction. Andriot and Fournier [4] used multiport systems and passivity to design control for teleoperated master-slave systems by H-infinity methods. Their strategy relies on the definition of a "distance to passivity" that measures the amount that a system must be changed to become passive. Along with this a "distance to transparency" is defined as the performance measure. These two measures are minimized by the H-infinity approach, with the passivity distance constrained to be negative, ensuring passivity.

Chapel and Su [19] have also treated the interaction of two port functions with robust stability methods, again based on passivity. Non-passive behavior of the robot port is modeled as a perturbation, and the small gain theorem is applied to ensure robust stability of the coupled system. This structure is used to derive conditions on the environment that guarantee stability with a certain robot controller, establishing coupled stability when the robot is nonpassive but "nearly passive." This has some striking similarities to the complementary stability approach proposed here. Complementary stability, however, more closely matches the physical structure of the problem. In practice, the behavior of the expected environments defines the problem, and conditions should be derived on the *robot* to satisfy stability with that set of environments. Chapel's approach does the reverse. Still, this is an effective formulation and could likely be substituted for the stability measure used here and produce similar results. Chapel does not propose a systematic design approach based on improving both stability and performance.

Colgate has proposed another method for establishing coupled stability with environments that does not require passivity [24]. Rather than base stability on the phase

criterion that flows from passivity, Colgate models the environment's port function as a perturbation on the robot impedance and derives stability by the small gain theorem, a magnitude criterion. Superficially this approach is even more conservative than passivity, but the author applies a series of coordinate transformations such that new, less conservative stability conditions can be derived for the robot based on the regions of the complex plane in which the environment's port function resides. This approach depends on finding a series of coordinate transformations that appropriately delineate the environment's properties; finding such transformations is the greatest challenge to applying this stability measure. To fully reduce conservatism, it is likely that these linear fractional transformations must vary with frequency. Colgate does not provide a systematic design procedure that could be used to provide improved performance while satisfying the stability constraint.

3.4.2 Measuring performance

In previous chapters, "feel" has been presented as a critical performance objective for human-interactive robots. Port behavior, in the form of mechanical impedance or admittance, provides a representation that captures the dynamic properties of the robot in isolation, as seen by the human interacting with it. Indeed, the desired feel can be quantified by a desired port function; performance error is deviation from this target function.

Newman has used high admittance as a performance measure for interacting systems. He has derived this from a desire for low contact forces when a manipulator moves along an interface that dictates position [87, 35]. Here the deviation from desired impedance is used. Impedance decreases monotonically, though not linearly, as admittance increases. Thus the two measures should produce similar, though possibly not identical results. Any difference could be compensated by frequency weighting.

In general control only has significant influence over the impedance in a finite range of frequencies. It is sufficient to evaluate performance at these frequencies, as higher frequencies do not yield any information to distinguish between solutions. Additionally, it may be useful to emphasize certain frequencies more than others.

This provides an opportunity for customization; for instance, a designer may wish especially to reject a certain low-frequency resonance, and therefore he may choose to weight more heavily the frequencies at which it dominates.

Practice dictates that deviation from the impedance magnitude should be considered on a log-log scale, as linear scales obscure the impact of low-frequency and low-magnitude effects. With logspaced frequencies, however, the impedance cannot be evaluated at DC. A lower bound on frequency must be chosen, along with a weighting function, that appropriately captures the importance of low-frequency performance without over-emphasizing it versus intermediate frequencies. The choice of frequency bounds and weighting function must be tailored to the application and robot system.

In general, the phase of the impedance is less intuitive than the magnitude, and poses several problems. One such problem is that two impedances with resonance at different frequencies differ dramatically in phase at frequencies between the two resonances. Another is that if desired behavior is passive, which it almost always is, constraining phase close to this target is very similar to imposing passivity, which is contrary to the goal of this approach. Finally, if the system is minimum phase, the phase and magnitude plots provide identical information, so there is no need to use the phase plot. In summary, there is no compelling reason to use phase, and there are several reasons to avoid its use, so only the magnitude is used. Should a situation arise where phase should be included in the performance measure, it can be added at that time.

Figure 3-9 shows the magnitude and phase of two typical systems. The solid trace shows the impedance of a two-mass, single-resonance system with large inertia (6 kg total) and damping (24 Ns/m total). The dashed trace shows the impedance of a single, lightly-damped mass (1 kg with 1 Ns/m damping), and might represent a target impedance. On the magnitude plot, the region between the two is shaded to illustrate the error. In the region below approximately 2 rad/sec, the difference between the shapes of the two curves results primarily from damping. At frequencies between 10 and 200 rad/sec, the slope of the two plots approaches unity and the difference in offset is due to the difference in apparent inertia. At higher frequencies, the resonance of the
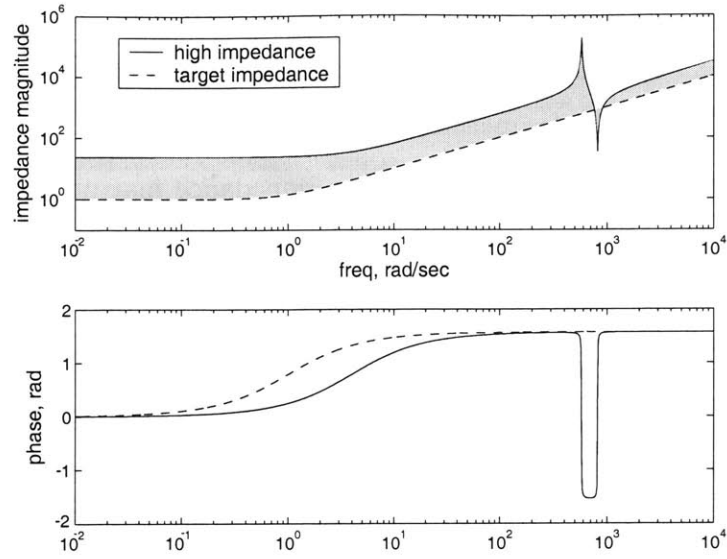
Figure 3-9: Magnitude and phase of port impedance for a lightly-damped mass (target behavior) and a single-resonance system with large mass and damping. The shaded region represents the error.

first system is apparent. At frequencies beyond this structural resonance, the offset between the two systems is reduced; this is because only half of the resonant system inertia (the portion closest to the interaction port) is apparent. In this region beyond resonance, a controller can have little or no effect on the endpoint because its action is attenuated by the intermediate dynamics. The phase plot shows the difference in the break frequencies of the two systems at low frequencies, and the resonance of the two-mass system around 600 rad/sec. If the system is controlled to reduce apparent impedance, the phase exceeds $\frac{\pi}{2}$ close to resonance (this is shown in the next chapter), and the phase error is likely to increase, even as the apparent friction and inertia are reduced. This is one reason that the phase is not generally used in the performance measure. The impedance magnitude offers an intuitive representation of the property to be regulated, and this is chosen as the performance measure.

The unweighted cost C_u is the deviation from desired impedance; that is, the shaded area in figure 3-9. The smaller the area, the closer to desired performance, and the smaller the unweighted cost. When the cost considers frequencies between

126

ω_0 and ω_1, a first pass at the unweighted cost may be:

$$C_{u1} = \int_{\omega_0}^{\omega_1} |(\log(|Z(j\omega)|) - \log(|Z_{des}(j\omega)|)|d\omega \tag{3.14}$$

(where to make the argument of the log dimensionless, both impedance magnitudes are normalized to 1 Ns/m). However, this arrangement for a continuous time measure weights the high frequencies much more heavily than the low frequencies. To determine the area as appears on a logarithmically scaled plot, as in figure 3-9, a change of variables is needed:

$$x = \log(\omega) \tag{3.15}$$

This leads to:

$$dx = d(\log(\omega)) \tag{3.16}$$

The "log-log" cost is then computed with the following expression:

$$C_u = \int_{x_0}^{x_1} |(\log(|Z(jx)|) - \log(|Z_{des}(jx)|)|dx \tag{3.17}$$

Here, $x_0 = \log(\omega_0)$ and $x_1 = \log(\omega_1)$. If a weighting function $W(\omega)$ is included, the weighted cost C_w is:

$$C_w = \int_{x_0}^{x_1} W(x)|(\log(|Z(jx)|) - \log(|Z_{des}(jx)|)|dx \tag{3.18}$$

As the controller design is done discretely on a computer, with a discrete frequency vector, the cost is approximated as a sum, rather than an integral:

$$C = \sum_{\omega_0}^{\omega_1} W(\omega)|(\log(|Z(j\omega)|) - \log(|Z_{des}(j\omega)|)| \tag{3.19}$$

When a logspaced frequency vector is used and $W(\omega) = 1$, this result is proportional to a discrete approximation of the area between the actual impedance Z and the target impedance Z_{des} as shown in figure 3-9. The logspaced frequency vector avoids the need for the change of variables.

127

3.5 Algorithm structure

The algorithm applied in this thesis proceeds through the following steps:

- A model is assumed for the robot, with parameters.

- A model is assumed for the environment, with nominal parameters and maximum uncertainty on each uncertain parameter.

- A performance metric is defined, in accordance with the guidelines provided in section 3.4.2.

- A form for the controller is assumed, with variable parameters P_i. If possible, this includes one parameter P_g (e.g. a gain) that is monotonically related to the cost and to stability, such that when all other control parameters are held constant, the limiting value of this parameter at which coupled stability is preserved provides the lowest cost. P_g should range over both stable and unstable values to find its stable limit. A parameter with these characteristics is not a requirement of the method, but it eases the interpretation of the results.

- The computation loop cycles through values for controller parameters P_i. P_g is varied last. For instance, assume the controller has three parameters, P_1, P_2, and P_g. For each combination of P_1 and P_2, the following procedure is implemented:

  1. P_g is incremented.

  2. Nominal coupled stability is checked by determining the location of the poles of the coupled system consisting of the robot and the nominal environment.

  3. If the system is nominally unstable, the algorithm goes to step 7.

  4. If the system is nominally stable, robust coupled stability is checked via computation of the singular value.

  5. If the system is robustly unstable, the algorithm goes to step 7.

128

6. If the system is robustly stable, the algorithm goes to step 1.

7. The cost for the previous increment of P_g is computed, and this cost and this P_g value are stored as the best achievable for this combination of P_1 and P_2.

- The grid of costs at the stability boundary is searched for the lowest cost. This is the best controller among those searched.

An example, with analysis and implementation of its results, is presented in the next chapter.

Chapter 4

Control Algorithm Validation

The control design procedure introduced in the preceding chapter is applied to an example using a linear model motivated by the literature. It is shown that the algorithm can converge to an improved, non-obvious controller. The resulting controller is implemented on the screw-driven single degree-of-freedom robotic system of examples 2.7 and 3.1, under the assumption that the linear model approximates the nonlinear robot system. System performance and stability are evaluated experimentally and compared to model predictions. Conclusions are drawn about the algorithm and the quality of the model.

4.1 Details of example

4.1.1 Robot model

A single-resonance robot model similar to the one used by Colgate [23] and Eppinger [38] and discussed in chapter 2 was used for interaction controller design, with parameters gathered from the characterization of the screw-driven vertical robot module [15]. An alternative model is a single inertia m with linear or nonlinear friction $F_f(x, \dot{x})$, driven by an actuator force F_a and environment force F_e, as shown in figure 2-24 and described by equation 2.18, which is reproduced here (a constant gravity force is

omitted for simplicity):

$$m\ddot{x} + F_f(x, \dot{x}) = F_a + F_e \qquad\qquad (4.1)$$

This model can be used to derive a simple impedance controller or a natural admittance controller for the system, as shown in chapter 2, but is not sufficient to predict the threat to coupled instability posed by force feedback because the (measured) environment force F_e and the actuator force F_a are collocated. An effective model must include, at a minimum, a single resonance. The real system in fact has many structural modes. This raises the question of how many modes should be modeled, a question that has been addressed by Colgate. In his analysis, Colgate found the same limiting gain in a distributed structural model as in a single-resonance model, suggesting that a single resonance captures the relevant non-collocated behavior [23]. In addition, the controller design procedure uses tools that require linear systems, so the nonlinear friction in the single-mass model cannot be used.

This method permits, in principle, the use of any robot model. Should it be necessary to use a high-order model, one can be inserted. For simplicity a single-resonance model, the minimum order model that demonstrates the relevant behavior, is used in this example. The model is shown in figure 4-1. This model is linear; the friction to ground is approximated with linear dampers b_1 and b_2. The total mass is represented by the sum of m_1 and m_2, whose motion is represented by positions x_1 and x_2, respectively, and their time derivatives. The stiffness k represents the structural resonance, and b_3 the structural damping. The parallel spring-damper k_f and b_f is a model of the force transducer, through which F_e is applied to the system. The actuator force is applied with F_a.

The robot and environment models are necessarily linear to apply the frequency-domain tools used here to evaluate both stability and performance. Coulomb and static friction must be omitted or approximated as linear damping. The implications of this are discussed below, in sections 4.3.5 and 4.3.6, in the context of the effect of various control strategies on port impedance, and in section 4.4.3, which analyzes data from robot implementation. Because this approximation is made, intuition pre-
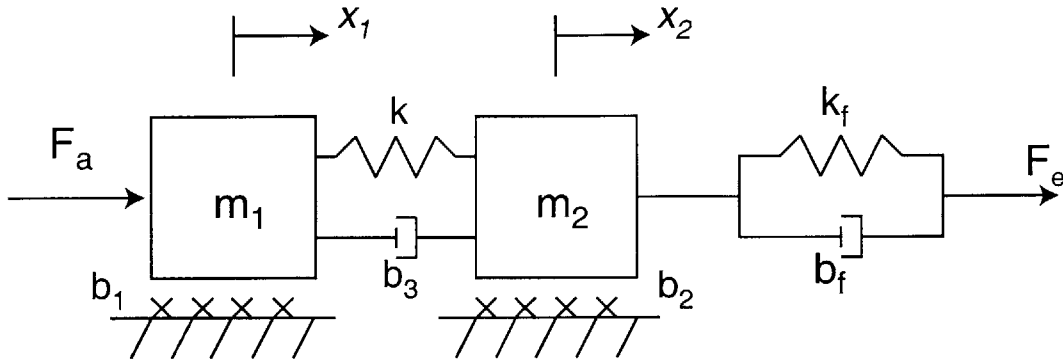
Figure 4-1: Single-resonance robot model with force transducer dynamics, for algorithm validation example.

dicts that the model-based method should become increasingly less accurate as static friction or other nonlinearities dominate intrinsic robot behavior. For instance, this approximation might be particularly poor when a system that is completely non-backdrivable, due to gear or screw binding, is controlled. However, this would be inconsistent with the guiding principle of this work: a combination of backdrivable hardware and control offer an approach to achieving interaction objectives; control alone can offer limited improvement, but cannot completely overwhelm the inherent characteristics of the robot [87, 15]. That said, the robot used for validation of this method has substantial Coulomb and static friction (as much as 20 N, 30% of the maximum force specification of 65 N), and offers a robust testbed for their approximation as linear damping.

The values used for the robot parameters are listed in table 4.1. The inertia and linear damping coefficients were determined from characterization experiments on the system, and symmetrically split between m_1 and m_2, and b_1 and b_2, respectively [15]. The force transducer is an ATI Gamma six-axis model; the stiffness for the z-axis is provided in the product documentation [7]. The damping coefficient was selected so that the resonance characterized by k_f and m_2 has damping ratio $\zeta \approx 0.05$, to prevent numeric problems that can result from undamped resonance. The structural stiffness k was selected based on separately-performed finite element analysis (provided by Interactive Motion Technologies) that predicted a first structural mode close to 90

Robot Parameter	Value	Units	Environment Parameter	Value	Units
m_1	3.0	kg	m_{hn}	2.1	kg
m_2	3.0	kg	m_{hd}	2.0	kg
b_1	12	Ns/m	b_{hn}	21	Ns/m
b_2	12	Ns/m	b_{hd}	20	Ns/m
b_3	1	Ns/m	k_{hn}	201	N/m
k	10^6	N/m	k_{hd}	200	N/m
k_f	17×10^6	N/m			
b_f	700	Ns/m			

Table 4.1: Robot and environment parameters for single-resonance robot example.
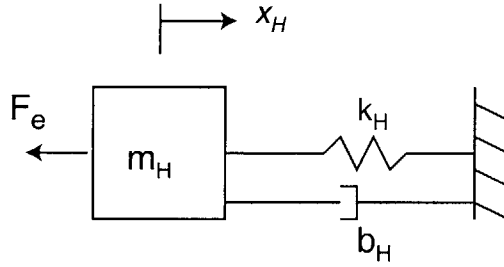
Figure 4-2: Second-order, single degree-of-freedom model of human arm.

Hz. The structural damping b_3 was made nonzero to avoid numeric problems, but kept very small in order to be as conservative as possible, as the exact value is not known.

4.1.2 Environment model

For the environment, a single degree-of-freedom linear model for a human arm was used. This model, shown in Figure 4-2, consists of a single mass m_h connected to ground through a spring k_h and damper b_h.

Each environment parameter is specified not as an exact value, but as a range of values, defined by a nominal and some maximum deviation from that nominal. Thus the actual range of possible parameter values is defined by the following set of equations, where m_{hn}, b_{hn}, and k_{hn} are nominal values, and m_{hd}, b_{hd}, and k_{hd} are the

maximum deviations for each:

$$
\begin{aligned}
m_h &= m_{hn} + \Delta_1 m_{hd} \\
b_h &= b_{hn} + \Delta_2 b_{hd} \\
k_h &= k_{hn} + \Delta_3 k_{hd}
\end{aligned}
\tag{4.2}
$$

Each Δ_i is a real number that can vary between -1 and +1:

$$
-1 < \Delta_i < 1, \Delta_i \in \Re, i = 1, 2, 3
\tag{4.3}
$$

The human musculoskeletal system is made up, in general, of rigid skeletal members, redundant muscles with nonlinear properties, and rotary joints, often with axes that move as the joint angle changes. It is controlled, with substantial time delay, via the brain and central nervous system, using sensory feedback from the eyes, muscles, and nerves in the skin. Despite the system's complexity, studies on fingers [47] and arms [92, 91] have shown local limb behavior to be dominantly low-order. The method presented here permits the use of any linear arm model; if a second-order model proves insufficient, a higher-order model can be substituted. This model uses a purely parametric representation of uncertainty. If arm behavior is dominantly second order, but can deviate from second-order behavior in certain frequency ranges, dynamic uncertainty could also be easily included. This would require a characterization of the uncertainty (including its magnitude versus frequency). Purely parametric uncertainties are used here to facilitate the verification and interpretation of results.

The values used for the stiffness, inertia, and damping parameters were selected to provide a representative range of typical human arm properties. There is a considerable literature devoted to the measurement of human arm impedance, particularly the stiffness, but most of this work focuses on planar arm movements, while the present application concerns vertical arm movements. Still, because approximate (order of magnitude) results are acceptable for this example, this literature provides a starting point for estimating vertical arm dynamics in a limited workspace in front of the torso. Simple measurements of the endpoint weight of an obese stroke patient's flaccid arm

suggested an effective endpoint arm mass of no more than around 2 kg. A similar result was found in [92], where a maximum of 3 kg across all subjects and conditions was found. Environment mass up to 4 kg was selected for controller design. Planar arm stiffness has been shown to vary with posture [50], with force level [92, 93], and to stabilize unstable environments [17]. Across a number of studies (e.g. [49, 50, 92]), single-joint stiffnesses were generally found to be less than around 50 Nm/rad. Assuming a length of 0.35 m of a single arm link (e.g. forearm), this is equivalent to an endpoint stiffness of around 400 N/m for small displacements. Other work that looks directly at endpoint planar arm stiffness (e.g. [17, 92, 93]) has found higher stiffness, as much as 1000 N/m or even 1500 N/m under certain conditions. These high stiffness values are usually obtained when a high endpoint force is applied, and when the endpoint stiffness is anisotropic due to kinematics. Assuming that a subject's arm is to be in a planar configuration, it is unlikely that vertical arm stiffness will reach these levels. As a first pass, the maximum environment stiffness was limited to 400 N/m. Should the result prove to be insufficiently conservative for interaction with human arms, the design procedure could be repeated with a higher stiffness value. Fewer studies have quantified the arm viscosity. In [56] an effective viscosity of 5.5 Ns/m was found. In [50], the joint viscosity was found to stay below 3 Nms/rad. Assuming a 0.35 m link length, this suggests an endpoint damping coefficient of around 25 Ns/m. Similarly, in [92] the maximum planar damping coefficient found across all subjects and all conditions was 40 Ns/m. This was used as the maximum damping value for controller design. The parameter values used for stiffness, damping, and inertia are shown in Table 4.1. Note that the undamped natural frequency ranges between 0.49 and 63.3 rad/sec. The impact on the results of changing the environment parameter ranges was not studied.

The form in equations 4.2 lends itself to representation of each arm parameter as a nominal with some bounded perturbation. If the robot is modeled as an admittance $\frac{1}{Z_{robot}}$ and the environment as an impedance, the coupled system is shown in Figure 4-3. The impedance function is broken up so that each parameter is independently perturbed. Coupled stability with all environments in the set is equivalent to robust
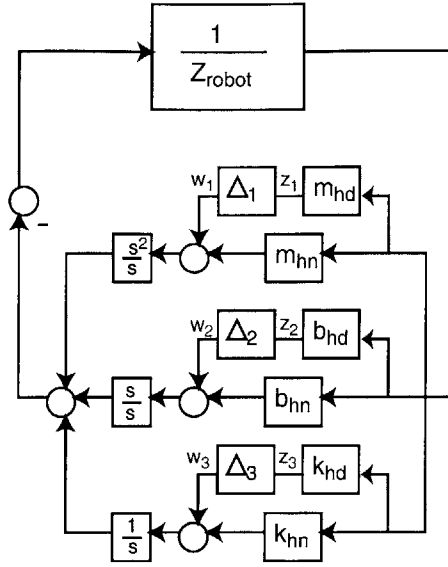
Figure 4-3: Block diagram of coupled system port functions. Environment properties are each perturbed.

stability of the loop in Figure 4-3.

4.1.3 Cost function and target impedance

For the initial algorithm testing, the target impedance has zero magnitude at all frequencies. This means that the ideal port behavior is that of a pure force source, perfectly backdriveable. This is an important behavior for stroke therapy robots, as they must avoid inhibiting patients' attempts to move during evaluation and for various times and directions of motion in the course of therapy. Also, active robot behavior is often so soft for this application that a zero impedance requirement is a good starting point.

The use of $|Z_{des}| = 0$ permits and requires simplification of equation 3.19. Because zero cannot be represented on a logscale, the desired impedance is removed from the expression altogether. Also, $|Z| \geq 0$ at all frequencies, and $W(\omega) = 1$ at all frequencies, so the cost between frequencies ω_0 and ω_1 can be computed as:

$$C = \sum_{\omega_0}^{\omega_1} \log(|Z(j\omega)|) \tag{4.4}$$

137

Because the log is taken of a quantity that might be less than unity, individual terms in sum 4.4, and therefore the sum itself, can be negative. Thus the cost has little absolute meaning. Its meaning can be captured by comparison between systems. For this example, ten thousand frequencies, evenly spaced on a logarithmic scale between $\omega_0 = .01$ and $\omega_1 = 1000$ rad/sec make up the frequency vector used to compute the sum C. The system under no control, $F_a = 0$, has $C = 19,720$. A pure inertia of 6 kg, the same as the system with no damping and no resonance, has $C = 12,782$. A 3 kg frictionless inertia, representing the theoretical limit of any passive controller, has $C = 9771$. These numbers provide a useful basis for comparison, but intuition is better served by examining plots of the impedance magnitude to understand the merits of specific controllers.

4.1.4   Controller form and robot port impedance

This method requires that the designer prescribe a structure for the controller and select one or more parameters for the algorithm to vary in search of the best control law. Any combination of motion and force variables can be fed back. Using the robot model in figure 4-1, a generic control law that feeds back the position of both robot masses, x_1 and x_2, and the measured interaction force F_e takes the Laplace domain form:

$$F_a = G_1(s)x_1 + G_2(s)x_2 + G_f(s)F_e \qquad (4.5)$$

In this control law, $G_1(s)$, $G_2(s)$, and $G_f(s)$ are all unspecified transfer functions. (To model velocity feedback instead of or in addition to position, it is sufficient to substitute sx_1 for x_1 and sx_2 for x_2.) Using this control law and solving the robot model equations of motion for the robot impedance (without force sensor) yields the

generic form:

$$Z_r(s) = \frac{Z_{rn}(s)}{Z_{rd}(s)} = \frac{m_1 m_2 s^4 + (m_2 b_3 + m_2 b_1 + m_1 b_3 + m_1 b_2)s^3}{m_1 s^3 + (b_3 + b_1)s^2 + ks - G_1 s + G_f s(b_3 s + k)} +$$
$$\frac{(b_2 b_3 + b_1 b_3 + b_1 b_2 + m_1 k + m_2 k)s^2}{m_1 s^3 + (b_3 + b_1)s^2 + ks - G_1 s + G_f s(b_3 s + k)} +$$
$$\frac{(kb_2 + kb_1)s - G_1 m_2 s^2 - G_1(b_3 + b_2)s}{m_1 s^3 + (b_3 + b_1)s^2 + ks - G_1 s + G_f s(b_3 s + k)} -$$
$$\frac{G_1 k + G_2 b_3 s + G_2 k}{m_1 s^3 + (b_3 + b_1)s^2 + ks - G_1 s + G_f s(b_3 s + k)}$$

(4.6)

When the force transducer stiffness and damping are included, the complete system port impedance is:

$$Z_p(s) = \frac{(b_f s + k_f)Z_{rn}(s)}{(b_f s + k_f)Z_{rd}(s) + sZ_{rn}(s)}$$

(4.7)

This is the port transfer function that is used to evaluate both performance and stability.

For the first example, only the force F_e is fed back, and a dynamic compensator $G_f(s)$ is designed to minimize the magnitude of the port impedance, $Z_p(s)$. This means that $G_1 = G_2 = 0$. The strategy is to begin with a simple structure, and to add complexity to search for improved performance. This was done by first assuming $G_f(s) = K_{dc}$, a constant gain proportional feedback law, and subsequently adding a pole and zero, first independently and then together.

Force control forms

Because the target impedance is zero, this example reduces to the special case of direct force control. Indeed, the port impedance and the stability and performance properties of the system do not change if a nonzero reference force F_{ref} is commanded and the compensator $G_f(s)$ acts on the force error $F_e - F_{ref}$. Many hundreds of papers have been published on force control, so it is worth considering common force controller forms in determining controller structures for this example.

As discussed in chapters 2 and 3, tools for the design of force and interaction controllers for stability and performance are not available at the same level as for servo controllers. As a result, many control structures have been proposed and used for force control, but their justification is mostly based on analysis of performance alone

139

or on empirical observation. Because investigators have used different assumptions and robot hardware, they have come to widely different conclusions about the relative merits of these forms. This is true even when discussion is restricted exclusively to the case of direct modulation of endpoint force (zero target impedance).

Proportional feedback of the endpoint force is one obvious strategy to regulate contact force. This strategy has the advantage of simplicity and has been deployed with modest success [15]. Its benefits can be seen using a simple robot model, as shown in example 2.7. Its simplicity has permitted its use in analyses of dynamic instability of force-controlled systems, for example in [38] and [23], and also as a baseline for comparison to alternative strategies, for example in [115] and [96].

Most literature on force control advocates some form of dynamic compensation. Several investigators have recommended including derivative action in addition to proportional feedback (PD control). One motivation is to improve stable interaction with stiff environments, as predicted by several models [39, 96]. Additionally derivative action is expected to improve the transient response to a change in commanded force. One major objection to implementing derivative force feedback is that computing derivatives of the measured force is prone to noise and time delay, particularly because force signals typically contain substantial high-frequency noise even before processing. It has been proposed to navigate this difficulty by implementing virtual damping via velocity feedback, which may provide many of the advantages of force differentiation [22].

Others have argued that in spite of its apparent benefits in improving coupled stability, derivative feedback is actually undesirable. Siciliano [107] and others make the case for proportional-integral (PI) action on the force in order to eliminate steady-state error. This is motivated solely by the performance objective of zero error. An argument for the stability of integral control action is provided in [116], but this uses a collocated robot model and makes the argument on the basis of delay in the control system. Still others argue for pure integral (I) control action [116, 120, 115], based on theory as well as experiments. Experimental evaluations of various force control strategies can be found in [22, 107, 115].

Additionally, model-based controllers have been studied and often used in conjunction with the control forms described above. Examples of this approach can be found in [107] and [84].

Some implementations use an inner motion-control loop in conjunction with a force control loop. This method is especially appealing for use with robots designed for motion control that have high intrinsic impedance and built-in low-level motion controllers. (Of course, this is contrary to the design philosophy espoused here, which is that hardware must be as backdrivable as possible so that the demands on force feedback control are tractable). In this approach, measured and desired force are compared and a desired motion trajectory is produced, which is then supplied to the low-level controller. Arguments for this include that motion control can achieve higher bandwidth through collocation, and therefore reject dynamic model errors and other force disturbances [32], and similarly that the inner servo loop can achieve high gains independent of desired output impedance, rejecting friction in the loop more effectively [107]. Natural admittance control [87], as described in chapter 2, is an example of this control style, and is studied further in this chapter.

Tuning rules are conspicuously absent from almost all of the force control literature. Papers that do suggest methods for tuning gains usually rely on explicit models of fully known environments, usually consisting of only a simple spring. The lack of tuning rules to place any controller poles and zeros, and the conflicting nature of the literature, makes the practical design of force controllers quite difficult. In motion servo control, it is generally considered bad practice to tune gains (e.g. PID gains) exclusively empirically, and the use of simple models and design tools such as root locus, Bode, or Nyquist design is preferred. In force and interaction control, these tools do not directly provide the necessary stability and performance measures. With the method proposed here, the goal is to provide a means of locating controller poles and zeros and tuning gains based on simple models and correct stability and performance measures, in order to move beyond empirical tuning as the prevalent method of controller implementation.

141

4.2 Computation

As this example is specified entirely by model, nominal stability can be determined by directly computing the location of the poles of the nominal coupled system. If the robot impedance is $Z_p = \frac{Z_{pn}}{Z_{pd}}$, and the nominal environment impedance is $Z_e = \frac{Z_{en}}{Z_{ed}}$, then the poles of the nominal coupled system are determined by finding the roots of the following equation:

$$Z_{pd}Z_{en} + Z_{pn}Z_{ed} = 0 \tag{4.8}$$

If any of the roots are in the open right-half plane (have real part greater than or equal to zero), the system is nominally unstable. If all of the roots are in the closed left-half plane (have real part strictly less than zero), the system is nominally stable.

Because Figure 4-3 has multiple perturbations Δ_i, it must be rearranged before the small gain theorem can be applied. The system must be put into the form of Figure 3-8, so all of the Δ_i terms must be put into a single matrix. This result is shown in Figure 4-4. The uncertainty matrix Δ is a diagonal matrix with Δ_1, Δ_2, and Δ_3 on the diagonal. The matrix M contains the rest of the system, here equal to:

$$M = \frac{\frac{1}{Z_{robot}}}{1 + \frac{m_{hn}s^2 + b_{hn}s + k_{hn}}{sZ_{robot}}} \begin{bmatrix} m_{hd}s & m_{hd} & \frac{m_{hd}}{s} \\ b_{hd}s & b_{hd} & \frac{b_{hd}}{s} \\ k_{hd}s & k_{hd} & \frac{k_{hd}}{s} \end{bmatrix} \tag{4.9}$$

The small gain theorem applied to Figure 4-4 via a singular value calculation produces a sufficient condition for stability. However, this condition is not necessary for robust stability in this case. The singular value approach is only necessary if all the terms in the matrix Δ can be varied to any value, real or complex, such that the infinity norm of Δ is less than one. The model used here has two types of structure that restrict Δ to a subset of the above: the off-diagonal entries are forced to zero, and the diagonal entries are real numbers. Thus application of the small gain theorem directly is conservative, and the structure of Δ dictates that the *structured singular value* be applied to find a necessary and sufficient stability condition.

Definition: The structured singular value μ of a complex matrix M with

142
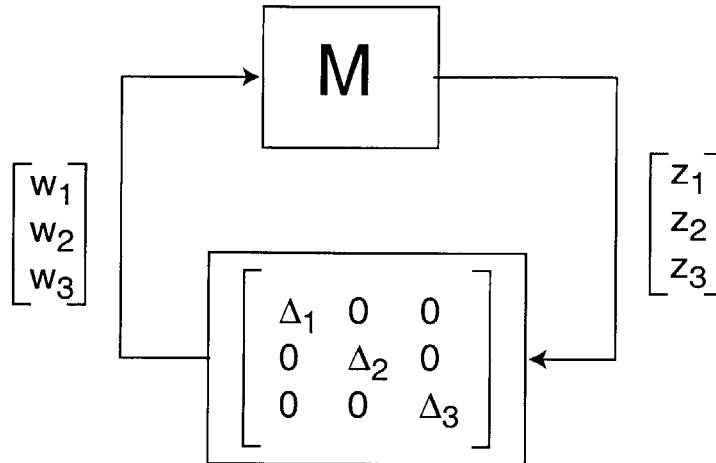
Figure 4-4: System with structured perturbation. The system in figure 4-3 can be rearranged to this form.

respect to a class of perturbations $\underline{\Delta}$ is given by:

$$\mu(M) \equiv \frac{1}{\inf[\sigma_{max}(\Delta)|\det(I - M\Delta) = 0]}, \Delta \in \underline{\Delta} \qquad (4.10)$$

If $\det(I - M\Delta) \neq 0$ for all $\Delta \in \underline{\Delta}$, then $\mu(M) = 0$ [30, 104].

In words, the structured singular value function finds the "smallest" admissible perturbation Δ (as measured by the peak singular value σ_{max}) that can make the characteristic equation $(I - M\Delta)$ singular. μ is the inverse of this smallest peak singular value. In this sense, it provides a measure of whether a set $\underline{\Delta}$ includes any perturbations that can make an $M - \Delta$ system unstable:

Theorem: An $M - \Delta$ system is stable for all $\Delta \in \underline{\Delta}$ with $||\Delta||_\infty \leq 1$ if and only if [30]:

$$\sup_\omega \mu(M(j\omega)) < 1 \qquad (4.11)$$

Large μ indicates that a small perturbation can make $(I - M\Delta)$ singular, and the system unstable. When $\underline{\Delta}$ is simply the set of all complex matrices, $\mu(M)$ is equal to the maximum singular value of M, such that the supremum of μ with respect to the

frequency ω equals the infinity norm of M:

$$\sup_{\omega} \mu = ||M||_{\infty}, \Delta \text{ arbitrary complex matrix} \tag{4.12}$$

The small gain theorem is a special case of the theorem provided above, for the case when Δ is an arbitrary complex matrix. The structured singular value provides a means to evaluate robust stability without conservatism when Δ has structure.

For some special Δ structures, μ can be computed exactly (for detailed properties of μ see [30, 104]). Unfortunately, when Δ is both real and diagonal, this is not the case. Instead, it must be estimated numerically.

4.2.1 Numeric issues

In general, there is no closed-form solution for the structured singular value problem. Some special cases can be solved analytically, but this is not true for this example, with Δ both diagonal and real. Structured singular values were computed with Matlab's μ-Analysis and Synthesis toolbox, particularly with the mu function (in this thesis, the mathematical function is denoted by μ, and the Matlab computation by mu) [8]. As with most numeric methods, this tool is not perfect, and several adjustments were made to ensure accurate computation. They are summarized below.

Pure real μ computation

Matlab computes the structured singular value by solving for an upper and lower bound at each point along a discrete frequency vector. When the uncertainty is purely real, as in this example, the bounds often do not converge. This results from the fact that when perturbations are purely real, there is no guarantee that μ is continuous with frequency. The problem is well known and is discussed at some length in the documentation for the toolbox [8]. The suggested fix is to double the size of both Δ and M. M_{new} is formed by simply appending a copy of M along the diagonal to the original. Similarly, Δ_{new} includes the original real perturbation structure $\Delta_R = \Delta$, and a copy Δ_C that has the same block-diagonal structure but that can have complex

144
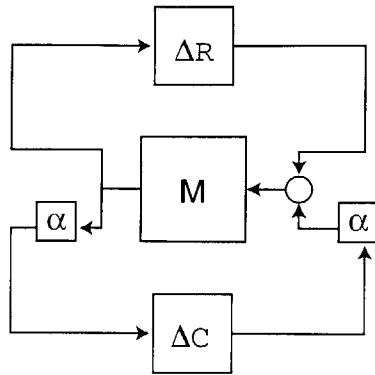
Figure 4-5: To ease μ computation in systems with purely real perturbations, a scaled complex perturbation that mirrors the real perturbation is recommended, as shown here [8].

values. The complex perturbations are scaled by some constant α^2 ($0 < \alpha < 1$):

$$M_{new} = \begin{bmatrix} M & 0 \\ 0 & M \end{bmatrix} \tag{4.13}$$

$$\Delta_{new} = \begin{bmatrix} \Delta_R & 0 \\ 0 & \alpha^2 \Delta_C \end{bmatrix} \tag{4.14}$$

This is depicted in Figure 4-5. The larger α is made, the greater the influence of the complex perturbation, and the greater the likelihood of convergence. Matlab provides both a physical and mathematical rationale for this fix [8].

Figure 4-6 shows an example of a mu calculation with Matlab's fix applied, with $\alpha = 0$ and with $\alpha = 0.5$. The lower bound computation fails completely in the case of a purely real perturbation. As α is increased, the bounds indeed do converge. Not surprisingly, the value of μ also increases substantially. This is because the fix introduces a broader set of perturbations than actually exist, providing more predicted opportunities for instability. The stability requirement is $\mu < 1$ for all frequencies; if the peak of the upper bound is less than one, the system is robustly stable. While the fix helps to validate the upper bound by making it converge to the lower bound, it also increases the upper bound, increasing its conservatism in comparison to the
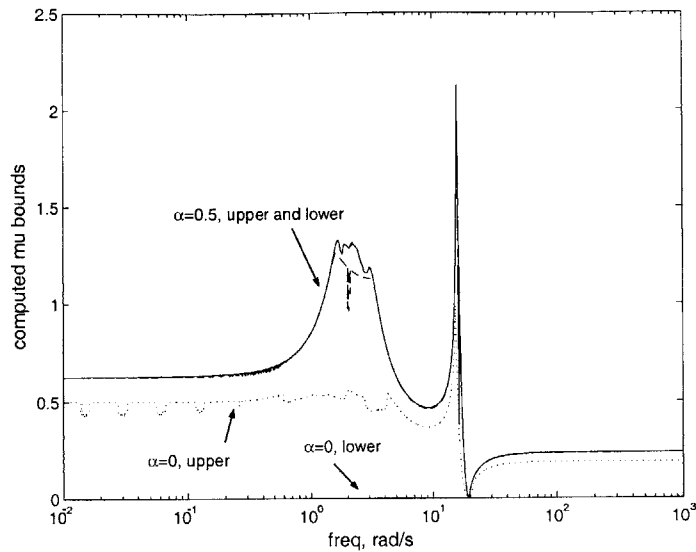
Figure 4-6: Upper and lower computed mu bounds with $\alpha = 0$ (purely real perturbation) and $\alpha = 0.5$. As α is increased, the upper and lower bounds converge, but both also increase.

actual real perturbation. Because only the peak of μ is relevant here, and because $\alpha = 0$ provides the least conservative estimate of the upper bound, the fix is not used for this example. The upper bound is by definition conservative, but it is not necessarily tight, because it is not guaranteed to converge with the lower bound. For this example, the lack of convergence is ignored, and to conserve computation resources, only the upper bound is computed. Because it is not known exactly how conservative the upper bound may be, it is always wise to perform some checks after the algorithm produces its results. One method of checking is described in the section below titled "Checking Results."

Frequency resolution

The upper bound is always conservative for the frequencies tested. Because a discrete vector of frequencies is used, however, the computed peak might differ from the actual peak of μ. Some small error is expected and is accommodated either by conservatism of the μ calculation, or by a stability margin introduced through use of a slightly lower gain. If the peak of μ is particularly sharp, however, and the spacing in the

146

frequency vector too sparse, larger errors could arise. The best way to avoid this is to densely populate the frequency vector. Unfortunately, this proportionally increases computation time.

To address this problem, a variety of methods could be applied to select the frequency vector to be locally dense at appropriate frequencies. One such method is described in the following section. Simple checking such as the approach described below should be used to catch significant errors and reduce the likelihood that an unstable solution escapes notice.

Avoiding unnecessary frequencies

Perhaps the most obvious approach to computing mu for a series of cases is to pre-define a frequency vector for all cases. However, there are at least two reasons not to do this. The first relates to the point made in the previous section. Frequency resolution must be high close to the peak of μ, but high resolution at all frequencies leads to long computation times and a wealth of unnecessary information. There is value in concentrating the frequency points close to the peak of μ, if possible.

A second, related point is that computing μ at unnecessary frequencies may produce misleading solutions. Because the upper bound of mu is not tight, it might in some cases predict instability at frequencies that can be shown by other means to be perfectly stable. In this example, the environment is passive, regardless of parameter values. At frequencies where the robot is passive, there is no risk of instability. The structured singular value need only be evaluated at frequencies where the robot is non-passive. With this approach, the only remaining source of conservatism is the lack of tightness in the upper bound of mu at those frequencies where stability is truly endangered, and unnecessary calculations are minimized.

A non-passive environment would not necessarily preclude an approach like this. It might be possible, for example, to determine the extrema of the phase excursion of the environment at each frequency. This could then be added to the robot phase, and if the total phase was between $-\pi$ and π, a sufficient condition for stability would be met. These approaches are not a good replacement for evaluating the

structured singular value, because they rely on sufficient but not necessary conditions, but they can be used to narrow the region in which μ must be determined, improving performance of the algorithm.

Checking results

A simple procedure was used to check some of the search algorithm results. Because the μ calculation relies on an upper bound, it is sufficient but not necessary for stability. A checking procedure was used that is necessary but not sufficient for stability; the actual stability boundary must be between the two, and if the results of both tests converge, the stability boundary is found. Minimum, nominal, and maximum parameters were selected for each of the three environment parameters (stiffness, inertia, and damping), and all 27 unique combinations of these parameters were determined. These models could then each be checked with a proposed controlled robot system model to verify stability or instability. If the μ-based algorithm predicts stability, all 27 environments should be stable when coupled to the robot. If the algorithm predicts instability, it is probable, though uncertain, that one of the 27 environments will couple unstably to the robot. This latter result is definitely not guaranteed, as the 27 selected environments represent only a small subset of the bounded infinite set of environments possible. However, because the subset includes all combinations of the extrema of the parameters (e.g. highest and lowest natural frequencies), it is likely that at least one member of this subset is among the most troublesome environments for a given controller.

If the search algorithm concludes that "controller A" should produce a robot that couples stably to all environments, and checking by this method indicates instability with one or more environment, this suggests a problem with the computation of mu, which should provide a sufficient test for coupled stability. One possibility is insufficient frequency resolution. If the search algorithm concludes that "controller B" should produce a robot that couples unstably to at least one environment, but checking finds stability with all environments, no certain conclusions can be drawn. It is possible that the mu calculation is conservative, or that a different environment

148

in the set, not one of the 27 tested, destabilizes the robot. If a slight increase in gain produces the expected instability, one of these two is likely the case, and the results should be accepted. If there is a large difference (this depends on the accuracy desired and the quantization of the gain vector, but results should be within perhaps 20%) in the gain at the stability boundary as determined by the two methods, further investigation should be pursued to resolve the discrepancy.

In over 90% of the cases tested in the example, this check agreed with the results of the mu calculation within the discretization of the gain vector used. In the few cases where the results of mu were more conservative than the check, the difference was less than 20% of the gain value. More importantly, these cases did not occur for controller parameter values that yielded good performance, so the difference had little impact on the results. This level of algorithm convergence is surprisingly good, considering the convergence problems of the real-valued structured singular value problem.

4.2.2 Computational Load

Because of the difficulty of computing structured singular values, the computational load associated with running the design algorithm is significant. On a personal computer with a 2.4 GHz Pentium 4 processor running Matlab Release 14 on a Windows XP platform, a calculation of mu using 1000 frequency points took approximately 20 seconds. Typical full results varied 3 parameters, usually along vectors 30 points long. By only calculating mu when absolutely necessary (when the robot was nonpassive but the nominal coupled system was stable) and by using previous results to localize the search over the gain parameter, the total typical search time was shortened to approximately 6 to 12 hours. Getting full results generally required several such runs, to adjust the parameter ranges to provide the most useful output.

Undoubtedly more streamlined search algorithms could be used to speed up this procedure. Despite the high computational load, it is important to note that this procedure is always done offline, to set gain conditions to be used in practice. Thus the high load is tolerable. As computing power continues to increase, problems with greater dimension can be managed.

4.3   Initial algorithm results

The algorithm was applied to the robot and environment model with several simple, classical control structures. In each case, the compensator was substituted for G_f in equation 4.7. Increasing the DC gain K_{dc} uniformly decreases the cost when all other parameters are fixed, so for each combination of the other parameters, K_{dc} was increased to its highest robustly (coupled) stable value, where the cost was then evaluated. Each vector of admissible parameter values can be any size; usually vectors of 30 were used for each parameter. In some cases multiple algorithm runs were used to cover additional frequency or gain ranges.

4.3.1   Proportional gain

The most basic controller tested uses constant-gain proportional force feedback:

$$F_a = K_{dc}F_e \qquad (4.15)$$

Because the gain K_{dc} is the only parameter varied, the computational load for this problem is quite low. The algorithm selects a maximum gain of 1.74; the resulting port impedance is plotted in figure 4-7. Because the gain exceeds 1, the system is expected to be nonpassive, and this is easily verified. While not surprising, this straightforward finding is significant; the complementary stability criterion, based on limited knowledge of the environment, admits a compensator that is invalid by the previous stability requirement (passivity). At least to the degree that the models of robot and environment reflect the physical systems they represent, this method locates acceptable controllers that outperform passive controllers. It is noteworthy that in practice significantly higher proportional gains (up to at least 5) have been used on the robot hardware represented in this simulation, without stability problems (see example 2.7). The algorithm appears to be conservative, at least for this case (though still less conservative than passivity). This is not entirely surprising; there are a number of potential discrepancies between the linearized models of both environment

150
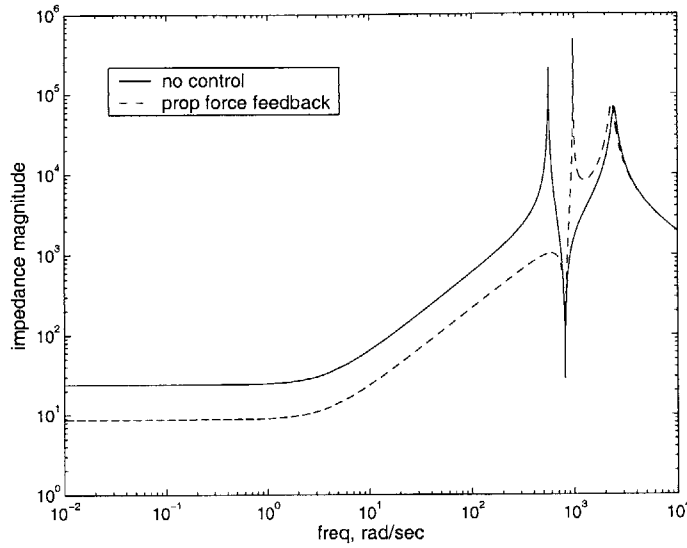
Figure 4-7: Port impedance magnitude of the robot system under no control, and under proportional force feedback with a gain of 1.74. Target behavior is zero impedance.

and robot and the physical hardware. Given these discrepancies, it is unrealistic to expect the algorithm to dictate numeric values for all gains that will translate exactly to the hardware. The real value of this algorithm is in rank-ordering controllers and suggesting what types of compensators are most effective. It is expected that if the algorithm predicts that controller A permits substantially higher DC gain than controller B, this will be borne out by the hardware. The algorithm should also provide order-of-magnitude estimates of stabilizing gains. Determining the exact value of the DC gain at which either controller is destabilized by a human arm, however, is not likely given the model simplifications used. The relationship between the model and the experimental results, and the conservatism of the model used here are discussed below, in section 4.4.2.

Given this, there is little value in the calculation of the expected maximum stable proportional gain, except to establish a baseline to which other controllers can be compared. With a proportional force feedback gain of 1.74, the system cost is 15,260. Only controllers that reduce the cost below this level are improvements on the obvious structure of proportional force feedback.

4.3.2 Single-zero (PD) control

Adding a zero with break frequency z (in radians per second) to the proportional controller of equation 4.15 produces:

$$F_a = K_{dc}\frac{1}{z}(s + z)F_e \qquad (4.16)$$

The factor $\frac{1}{z}$ is included so that K_{dc} indicates the gain at DC. If the control law is expanded and rewritten in the time domain:

$$F_a = \frac{K_{dc}}{z}\dot{F}_e + K_{dc}F_e \qquad (4.17)$$

it is evident that it is equivalent to proportional-derivative (PD) control. While such a controller cannot be implemented in practice, it can be approximated by adding a high-frequency pole, and it is useful to explore whether derivative action improves performance.

The zero break frequency z was varied along discrete intervals from 0.01 to 10^6 rad/sec. Although it is not practical to implement controllers at such high frequencies, it is useful to consider the idealized results of this control. At each frequency, the gain K_{dc} was elevated to the highest level that preserved coupled stability with all environments. The cost at the maximum stable gain is plotted versus zero frequency in figure 4-8. With a zero at low frequencies, the cost is higher than with proportional control. This suggests that pure derivative feedback has little value. At very high frequencies, performance approaches that of proportional control, as expected. When the controller zero is sufficiently removed from the system dynamics, it has little effect. With a controller zero around 2000-3000 rad/sec, the cost is around 20% lower than it is with proportional control. This indicates that placing a control zero at frequencies just above the robot's resonant frequency (around 500 rad/sec) may have some value. Other single-zero compensators are ineffective in improving performance versus proportional force feedback.

152
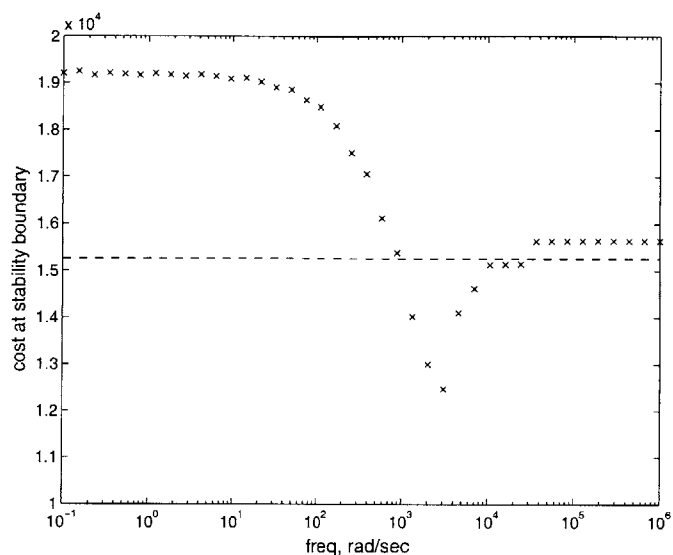
Figure 4-8: Cost at the maximum stable gain versus zero frequency. Single-zero force compensators. Dashed line shows the cost at the limiting proportional feedback limit.

4.3.3 Single-pole (lowpass) control

Adding a pole with break frequency p (in radians per second) to the proportional controller of equation 4.15 produces:

$$F_a = K_{dc}p\frac{1}{(s+p)}F_e \qquad (4.18)$$

This controller has lowpass properties. p was varied discretely from 0.01 to 10^6 rad/sec. At each frequency, the gain K_{dc} was elevated to the highest level that preserved coupled stability with all environments in the expected set, and the cost was evaluated. The cost is plotted versus the pole break frequency p in figure 4-9. With p at high frequencies, performance approaches that of the proportional controller. Placing a controller pole at moderate frequencies roughly centered around the robot's 500 rad/sec resonance degrades performance significantly. Performance is best when the pole tends toward very low frequencies. From figure 4-9 it appears that moving the pole to even lower frequencies might improve performance further. This suggests that pure integral force control may be an effective way to reduce apparent

153
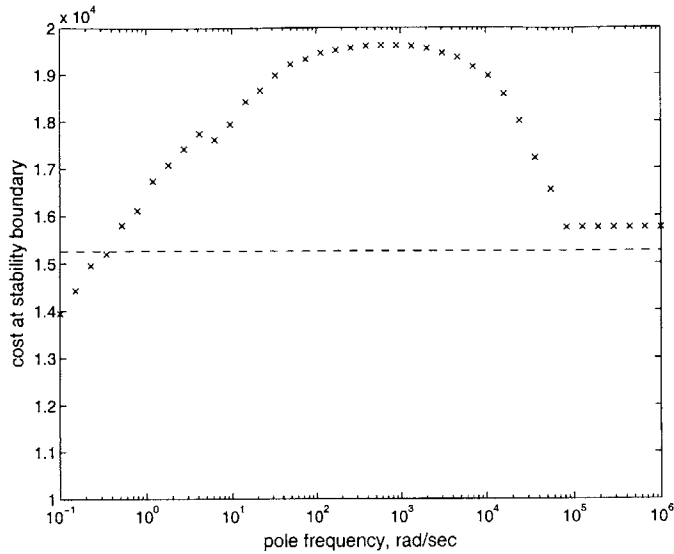
Figure 4-9: Cost at the maximum stable gain versus pole frequency. Single-pole force compensators. Dashed line shows the cost at the limiting proportional feedback limit.

endpoint impedance. This idea is further discussed in the analysis of lag controllers below.

4.3.4 Lead and lag controllers

Another simple classical control form that offers greater flexibility includes a single pole at frequency p as well as a single zero at frequency z and has the form:

$$F_a = K_{dc} \frac{p\,(s+z)}{z\,(s+p)} F_e \tag{4.19}$$

This form admits lead compensators, when $z < p$, and lag compensators, when $p < z$. When $z = p$ the pole and zero cancel and the form reduces to proportional control.

z and p were each assigned identical discrete frequency vectors. At every unique combination of pole and zero frequencies, the gain K_{dc} was raised to its maximum stable value. The cost at the stability boundary is plotted versus pole and zero frequency in figure 4-10. Pole frequency p is along the x-axis, zero frequency z along the y-axis, and the cost is plotted along the z-axis.

154
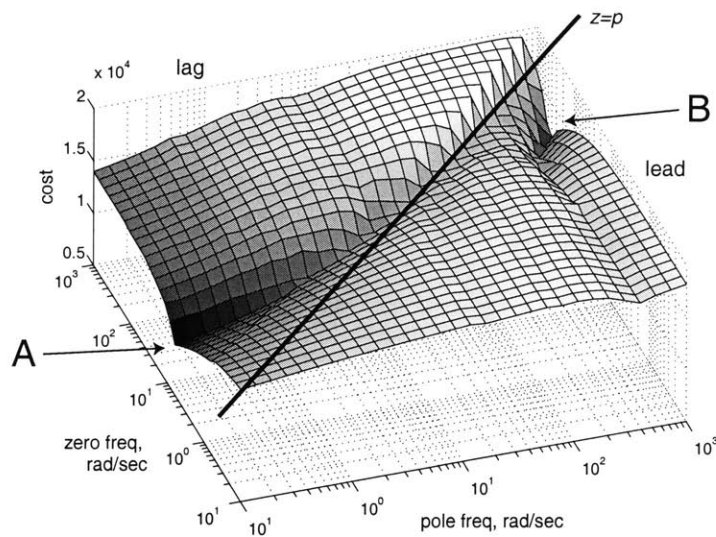
Figure 4-10: Cost at the stability boundary for controllers with a single pole and single zero, versus pole and zero frequencies.

The diagonal $z = p$ shows the same cost for all z and p values, as expected for proportional control. Above and to the left of this diagonal, $p < z$ which indicates lag compensation. Below and to the right of this diagonal, $z < p$ indicating lead compensation. There are two promising regions, where the cost is significantly lower than for proportional control. The first is in the lag region, with poles at the lowest frequencies and zeros at intermediate frequencies, with the lowest cost around $z = 6$ rad/sec. In this region, marked "A" in figure 4-10, the cost is as low as 9800 on the plot, but is even lower when the pole is moved to lower frequencies. With $p = 0.01$ rad/sec and $z = 5.7$ rad/sec, the cost is 8977. The second promising region, marked "B" in the figure, is in the lead region, at high frequencies. The lowest cost evaluated in this region is 11,640 with $z = 281$ rad/sec and $p = 728$ rad/sec.

Figure 4-10 provides an opportunity to compare lowpass, single-pole control to lag control. For any pole frequency below 1 rad/sec, performance is always superior if a zero is included around 6 rad/sec versus a zero at high frequencies. For $p = 0.1$ rad/sec, from figure 4-9 the cost is 14,000 if there is no controller zero. From figure 4-10, for $p = 0.1$ rad/sec and $z = 5.7$ rad/sec, the cost is 9800. So, despite the promise of

155

low-frequency poles without zeros (approaching pure integral control as $p \to 0$) shown in figure 4-9, low-frequency lag control (approaching proportional-integral control as $p \to 0$) appears more promising.

4.3.5 Results analysis

Lag compensation

Looking at the impedance plots with both controllers shows the difference between the lag and lead approaches, and illustrates their competing features more clearly than the cost. Figure 4-11 is a magnitude plot of the port impedance of the system without control, with proportional control at the stable gain limit, and with lag control at the stable gain limit. The apparent inertia, indicated by the section between 6 and 300 rad/sec, is nearly identical between the proportional and lag controllers (any difference is likely a result of the discretization of the DC gain vector). All of the benefits from the lag compensator (versus proportional control) arise at low frequencies, while the impedance under lag control converges almost identically to the impedance under proportional control at higher frequencies. The loss of passivity under force feedback is due to the structural resonance, and the coupled stability problems generally arise in frequencies surrounding this resonance. The lag controller is sculpted to boost the gain at low frequencies but to avoid changing anything at the frequencies where coupled stability is at risk.

The selection of the zero break frequency for the lag compensator is critical, as is shown by the dramatic improvement in cost around $z = 6$ rad/sec in figure 4-10. The reason for this can be seen by looking at the phase plot of the system impedance with several different lag compensators, shown in figure 4-12. The solid trace is the system without control; note that this is the only passive trace plotted, as its phase stays between $-\frac{\pi}{2}$ and $+\frac{\pi}{2}$ for all frequencies. Proportional control and several lag compensators are also shown; each of these is nonpassive around 700 rad/sec, because of the non-colocated force feedback. The lag control produces a bulge in the phase at low frequencies. If the zero break frequency is significantly
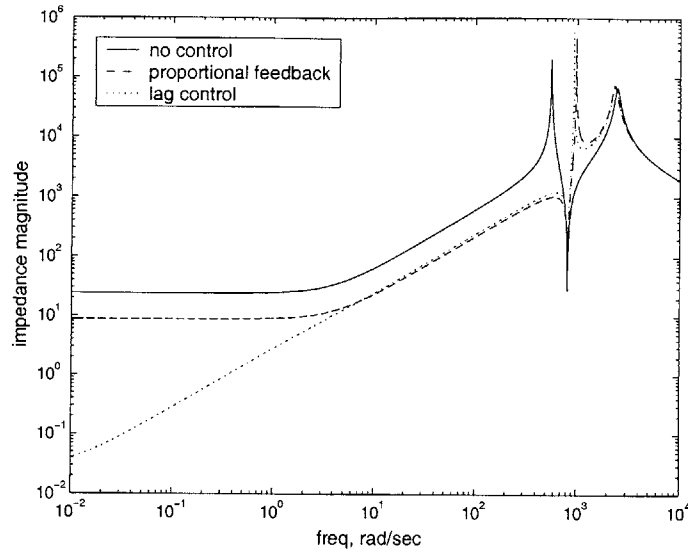
Figure 4-11: Port impedance magnitude of robot with no control, proportional control, and lag control.

higher than 6 rad/sec (e.g. 20 rad/sec, which is plotted), the bulge exceeds $\frac{\pi}{2}$, and the robot system is nonpassive at low frequencies as well as around the structural resonance. If the zero frequency is too low, it moves closer to the pole and reduces the benefits of lag control at low frequencies. The algorithm identifies the appropriate break frequency to keep the system passive at low frequencies, but to take advantage of lag action. At high frequencies, the effective lag compensator has phase nearly identical to the proportional controller. This is because both controllers have the same characteristics at these frequencies, and both encounter the same limits to stability. The lag compensator simply boosts the gain at frequencies significantly removed from the structural resonance. It is particularly good at rejecting low-frequency friction.

At frequencies above resonance, some compensators with sufficiently high gains are nonminimum phase. This is discussed in appendix A.

Lead compensation

The algorithm predicts that lead compensation at high frequencies is also beneficial in reducing cost, but analysis of the port impedance under this control reveals that
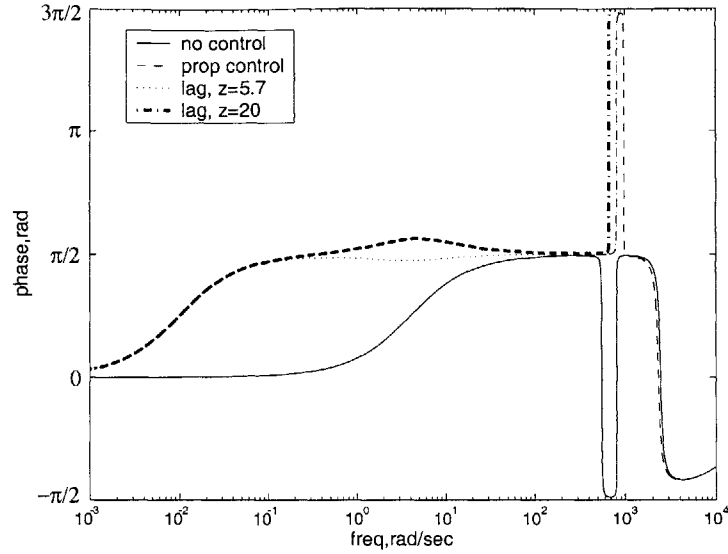
Figure 4-12: Robot port impedance phase under no control, proportional control, and two different lag controllers.

the mechanism by which it helps is quite different than lag control. Figure 4-13 shows the magnitude of the port impedance of the system under no control, at the proportional stability limit, and with high-frequency lead control at its stability limit. Lead compensation reduces the apparent impedance roughly equally at all frequencies below approximately 200 rad/sec. The apparent inertia is reduced substantially below the other control methods; friction is also reduced, but not nearly as effectively as with lag control. A phase plot of the robot port impedance with no control, proportional control, and lead control is shown in figure 4-14. At frequencies close to the structural resonance, the selected lead compensator reduces the phase excursion of the port impedance, increasing the lowest frequency at which the robot port is nonpassive, and reducing the amount that the phase exceeds $\frac{\pi}{2}$. The algorithm selects a controller such that this happens at the correct frequency range to permit the DC gain to be raised substantially. This is less intuitive than the effect of lag compensation; the high-frequency lead adjusts the port impedance phase at high frequencies to stabilize the coupled system at higher gains.
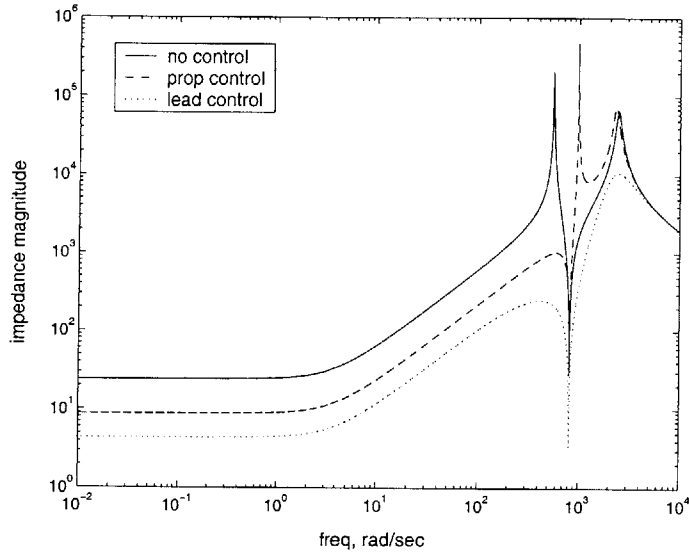
Figure 4-13: Port impedance magnitude of system under no control, proportional control, and lead control.
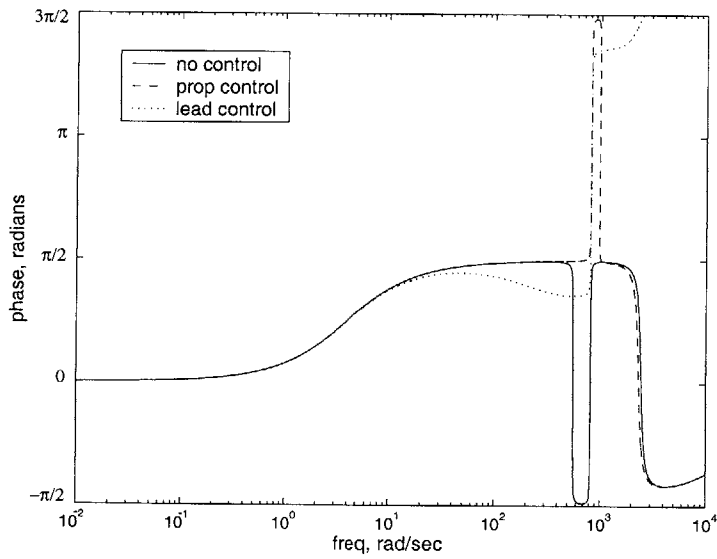
Figure 4-14: Port impedance phase of system under no control, proportional control, and lead control.

159

Lag-lead compensation

Because low-frequency lag and high-frequency lead compensation offer different means of reducing the apparent impedance magnitude, combining the two in some way seems promising. Because the initial search only included compensators with a single pole and a single zero, it is not capable of identifying lag-lead or lead-lag controllers. Including an additional pole and zero in a new search is certainly feasible, but it would increase the computational load substantially by going from three dimensions to five.

The zero break frequencies of the lag and lead compensators are at 6 and 281 rad/sec respectively, more than a decade apart. This suggests that the two pole-zero pairs might be superposed without interfering with each other's benefits. Including two pairs of poles and zeros in a single control law produces:

$$F_a = K_{dc} \frac{p_1 p_2}{z_1 z_2} \frac{(s + z_1)(s + z_2)}{(s + p_1)(s + p_2)} \tag{4.20}$$

In this control law, p_1 and p_2 are the pole break frequencies, and z_1 and z_2 are the zero break frequencies. Using information from the lag and lead compensators found in the previous controller search, the values of all four break frequencies were fixed to provide low-frequency lag ($p_1 = 0.01$, $z_1 = 5.7$) and high-frequency lead ($z_2 = 281$, $p_2 = 728$). With these parameters fixed, the gain K_{dc} was varied to find its maximum value that preserves coupled stability. The resulting port impedance magnitude can be seen in figure 4-15. This compensator captures some of the benefits of the lag compensator and the lead compensator described above, as it substantially reduces low-frequency friction as well as the inertia. The magnitude and phase of the three controllers as well as the proportional controller at the stability limit are plotted together in figure 4-16. The lag-lead controller phase tracks the lag controller's at low frequencies, and the lead controller's at high frequencies. The high-frequency lead allows higher controller magnitudes at low frequencies and across the frequency spectrum, as compared to the lag compensator. The cost associated with this controller is 7020.

While the form of each of the controllers listed above is relatively simple, the key

160
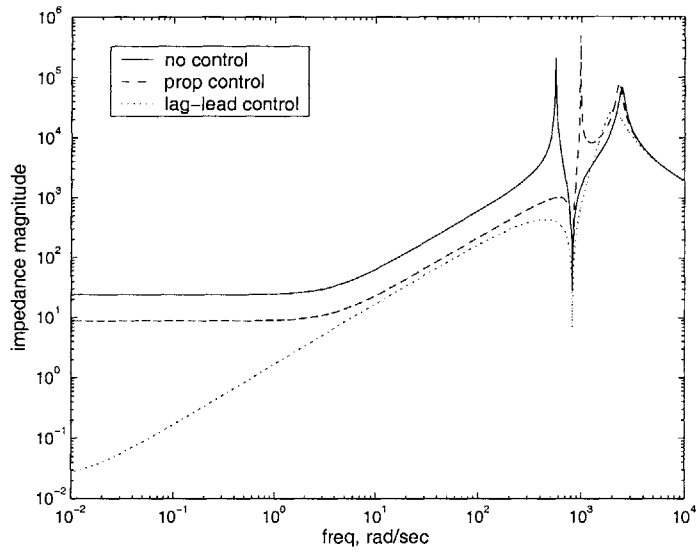
Figure 4-15: Port impedance magnitude of system under no control, proportional control, and lag-lead control.
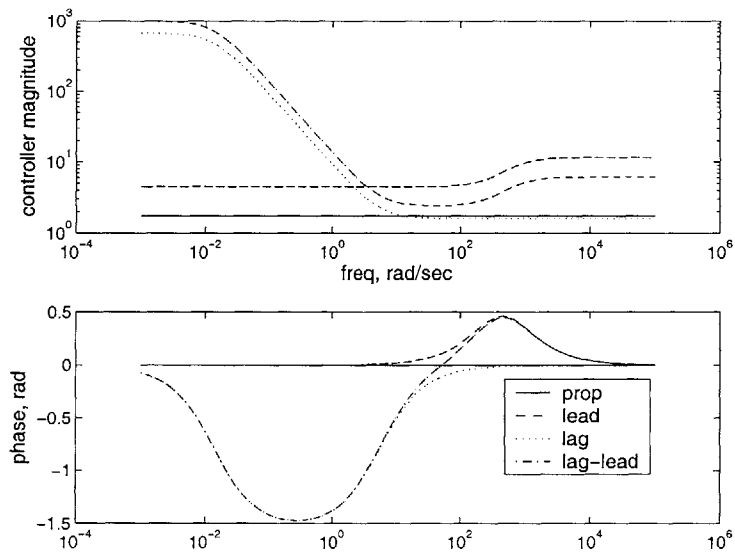
Figure 4-16: Controller magnitude and phase for proportional control, lag control, lead control, and lag-lead control.

point is that the algorithm provides guidance in each of their designs. The results are not intuitive and could not be arrived at by applying classical control design tools (like root locus, Bode/Nyquist design, etc.). Without actually simulating particular control forms and parameters, the impact that control has on the impedance phase, as shown in figures 4-12 and 4-14, is not obvious, and is not trivial to determine. Additionally, even if one determines the effect that a lag compensator, for example, has on the impedance phase, it remains non-obvious how to change the compensator to affect desired changes to the phase. This results directly from the structural differences between servo and interaction control design as described in section 3.3.1, and the results of this example verify the problems identified therein. The algorithm successively makes changes to the compensator and evaluates the effects of these changes against an objective standard, providing a systematic method to design interaction controllers. The algorithm also identifies where and how the system can relax passivity, becoming more non-passive in ways that improve performance without compromising coupled stability. Some insight can be gained after the fact by examining the features of the most promising controllers, as in this section, but because of the non-intuitive way that controller features impact port impedance, this is not enough to enable design by a traditional method.

Influence of environment damping

Because coupled instability results from the generation of energy by one or more of the coupled systems, one common perception is that increasing the level of physical damping in the system can drain sufficient energy to stabilize an otherwise unstable connection. In this section it is shown that increasing damping can help or hurt stability, and that this depends very specifically on the properties of the robot and the nature of its controlled port impedance.

The environment is modeled as a mass with a spring and damper connecting it to ground. Since the robot has been characterized as an impedance, the environment is

162
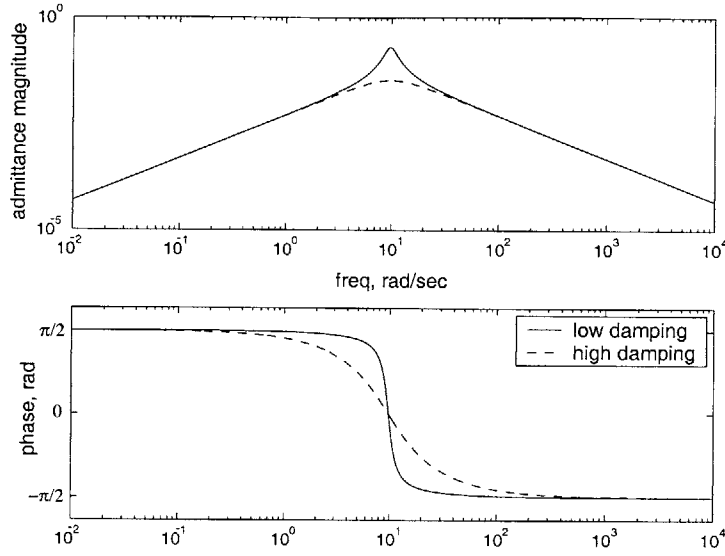
Figure 4-17: Environment admittance, magnitude and phase.

modeled as an admittance described by the equation:

$$Y_{env}(s) = \frac{s}{m_h s^2 + b_h s + k_h} \tag{4.21}$$

The magnitude and phase of this admittance for two separate values of b_h (and fixed values of m_h and k_h) are plotted in figure 4-17. The phase of this admittance always starts at $+\frac{\pi}{2}$ at DC and tends toward $-\frac{\pi}{2}$ at high frequencies, with a second-order break point due to the resonant dynamics. When b_h is small, the phase breaks sharply. As b_h is increased, this break becomes more gradual, and the phase is decreased at frequencies below resonance and increased at frequencies above resonance.

From figures 4-12 and 4-14 it is clear that when passivity is lost due to non-collocated force feedback, the robot port impedance phase exceeds $+\frac{\pi}{2}$, and that the phase approaches $\frac{3\pi}{2}$ at its maximum (except for the nonminimum phase behavior discussed in appendix A). If the product of the robot port impedance and environment admittance $Z_p(s)Y_{env}(s)$ is the open-loop transfer function of the unity negative feedback system that represents the coupled system, then the phase of the two port functions adds to give the open-loop phase. Coupled stability is risked when

163

this total phase exceeds π. Adding damping to the environment decreases phase in some regions but increases it in others. Thus the relationship between the environment's resonant frequency and the frequencies at which the robot is nonpassive (which can vary substantially on the basis of control properties as well as physical robot properties) determines whether damping in the environment has a stabilizing or destabilizing influence. In fact, with certain control parametrizations tested here, the most destabilizing environments have the maximum damping, and with others, the most destabilizing environments have the minimum damping. The role of environment damping in helping both to stabilize and to destabilize contact is shown by way of example in section 4.4.2.

It is likely true that modifying damping somewhere in the system could almost always improve stability properties (Dohring shows one example of strategically placed damping, which is described in the next chapter [35]). However, the precise placement of this damping is critical, and depending on the situation, it may be better to increase or decrease its magnitude.

4.3.6   Nonzero target impedance

The example described to this point pursued the objective of zero robot port impedance, or force control. It is required that the system also be capable of representing nonzero impedance, particularly that it render virtual spring-damper combinations of a range of values. The control design procedure provides the means to include this by adjusting the cost and by changing the form of the controllers.

One way to introduce nonzero impedance behavior is to include position and velocity feedback, as well as force feedback, each with dynamic compensators, and to search across all combinations of parameters for those compensators to find a controller that best provides the virtual impedance requested while minimizing all other unwanted impedance properties. This would require a massive search over many parameters, as each controller would have to be of sufficiently high order to provide both the desired behavior and some freedom to suppress undesired properties. This is certainly possible within the framework described here, but is computationally

164

burdensome.

Instead, since virtual spring-dampers are specifically targeted, and because a ready means of producing these characteristics with motion feedback is available, an alternative approach is taken. Constant-gain position feedback is used to create stiffness, and constant-gain velocity feedback is used to create damping (as in simple impedance control; see chapter 2). The position and velocity gains are preset to the desired stiffness and damping, respectively. (Because the viscous damping in the robot model in part approximates unwanted nonlinear friction, it is not used to provide desired linear damping; instead the strategy is to introduce desired damping with control, and suppress all physical friction with force feedback.) A dynamic force compensator is used to minimize the force error, defined as the deviation of the measured force from that specified by the impedance law. The control law takes a form very similar to equation 2.17, except that the proportional force feedback gain K_f is replaced with a dynamic compensator. The design algorithm searches among the parameters of this compensator for the best solution. Assuming the virtual trajectory $\dot{r} = 0$ and $r = 0$, and assuming a single-zero, single-pole form for the force compensator, the control law takes the form:

$$F_a = -Kx_1 - B\dot{x}_1 + K_{dc}\frac{p\,(s+z)}{z\,(s+p)}(F_e - Kx_1 - B\dot{x}_1) \qquad (4.22)$$

Although the transfer function by which the measured force F_e is multiplied is exactly the same in this case as in the zero impedance case, the rest of the control law changes the properties of the physical robot system. Running this algorithm tests whether these changes substantially impact the optimal form of the force compensator.

Nonzero stiffness, zero damping

Adding stiffness to the system has a dramatic impact on its impedance structure, as behavior (below structural resonance) is changed from that of a damped inertia to that of a spring-mass-damper. Low-frequency behavior is dominated by the stiffness, high-frequency by the inertia, and damping mainly affects the nature of the primary
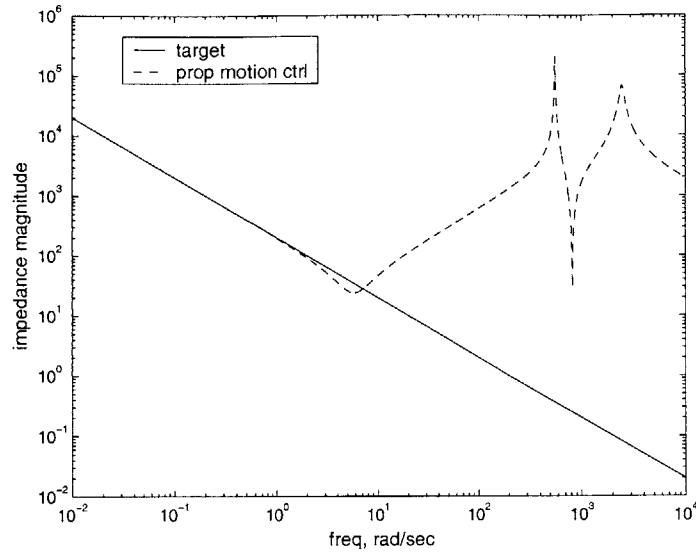
Figure 4-18: Robot port impedance magnitude under proportional control, with a target stiffness of 200 N/m.

resonance. A magnitude impedance plot of the system with proportional control only (no force feedback) and a desired stiffness of 200 N/m is shown in figure 4-18. At low frequencies there is little or no difference between the target behavior and the actual system behavior. The motion feedback controller adequately replicates spring behavior at low frequencies. The damping (approximating friction) and inertia only cause deviation from the ideal at moderate to high frequencies. This is not true of nonlinear friction like Coulomb and static friction, which have significant low-frequency and DC components. Thus the use of linear damping to approximate friction is less effective here than in the zero-impedance case.

The search algorithm was repeated for a series of cases with zero target damping and various target stiffnesses, including $K = 200, 2000$, and 20000 N/m. In general, for each pair of pole and zero frequency parameters, the maximum stable gain is very similar for these cases as for the zero impedance case; as the target stiffness gets very high, lower force feedback gains are required for some controllers. For instance, with virtual stiffness $K = 20000$ N/m, DC feedback gains must be lowered by between 0 and 60% from their levels at zero target impedance. The high-performing lag and

lead compensators, however, can achieve gains as high (or very nearly as high) as the zero-impedance case. The rank-ordering of controllers, determined by comparing controllers with different pole and zero frequencies to one another, is preserved with respect to the maximum stable gain. Low-frequency lag and high-frequency lead compensators remain the best performers.

When the cost is considered, however, there is one substantial difference between the zero and nonzero stiffness cases. Because they act only to increase the gain at low frequencies, lag compensators show virtually no reduction in cost as compared with proportional force compensators. Because the motion feedback controller does an adequate job of rendering low-frequency stiffness, there is little room for improvement by locally boosting the force gain at low frequencies. Of course, because in the real system the nonlinear friction actually does have a strong impact at low frequencies, there is a benefit to these controllers, but this benefit is not identified by the method. There are several ways to remedy this problem; for example, a measure of force gain, at all frequencies or in a particular band, could be explicitly included in the cost. The wiser choice may be, however, to focus controller design on the zero impedance case, and to only tweak it if necessary for nonzero cases. This approach is further justified below.

Zero stiffness, nonzero damping

If the stiffness $K = 0$ but the damping B is positive, a slightly different condition arises. In this case the system impedance more closely resembles that of the system with zero target impedance.

Several runs of the control design algorithm (with $B = 10$ and $B = 50$ Ns/m) predict that as the desired damping increases, the force feedback gain can be raised substantially without compromising stability. For example, if the virtual damping is set at $B = 50$ Ns/m (a value equal to more than double the intrinsic system damping), the predicted stable DC gain limits increase by more than 400% for most pole-zero combinations (including the high-performing ones that are likely to be selected), and by as much as two orders of magnitude in some cases. One notable

167

exception is the high-frequency lead compensators, whose limiting gain stays approximately constant. Thus these compensators are not favored over proportional control when virtual damping is added. The designer should be somewhat skeptical of this outcome, because the idealized model assumes that damping implemented with the controller is identical to physical damping in the system, which is not necessarily the case. The algorithm may overestimate the stabilizing properties of the active damping. At the very least, though, one can note that a controller designed for the zero impedance case should satisfactorily stabilize interaction when nonzero damping is specified. Low-frequency lag compensators again achieve the highest gains. However, as the virtual damping is increased, the optimal break frequency of the lag compensator's zero shifts gradually upward. With $B = 50$ Ns/m, the highest gains are achieved when the zero is between 50 and 80 rad/sec. If the zero is placed at 6 rad/sec as selected by the zero impedance design, however, gains can still be achieved that are at least five times greater than in the zero impedance case. It is noteworthy that this is consistent with much of the force control literature, that suggests adding substantial virtual damping to improve stable contact [22, 32].

Nonzero stiffness and damping

No surprises emerge when nonzero virtual stiffness and damping are proposed; the predicted gain limits are between the prediction for stiffness alone and damping alone. For example, if $K = 200$ N/m and $B = 10$ Ns/m, the limiting gains are greater than for the case when $K = 200$ and $B = 0$, and less than for the case when $K = 0$ and $B = 10$.

Overall, nonzero virtual stiffness and damping make little difference in the selection of optimal pole and zero locations, yielding results that are similar, if not identical, to the analysis done for zero target impedance. This implies that, under the assumptions inherent to this model, zero-impedance force control is the critical objective, and that achieving nonzero impedance requires only relatively small modifications. However, it is important not to extrapolate this conclusion beyond the assumptions. The model used here possesses several features that might lead to this

result. The motion sensors and actuator are assumed perfect and collocated, so that the system is ideally suited to implement desired stiffness and damping independent of any force-based control strategy. The robot is stiff and therefore effective at transmitting the ideally-implemented actuator behavior through to the interaction port. In other words, the modeled system is suited to accurately produce nonzero port impedance, aside from the high inertia and damping. Finally, the impedance control law is structured so that the controller minimizes force error, and is very similar whether target impedance is zero or nonzero. For this model, achieving adequate force control is the critical performance objective that leads to performance in all specified conditions. This principle should be used to guide design of a controller for this model, but not necessarily for design of hardware. For different physical systems, the needs may be quite different.

Designing for a range of impedance

In practice, a human-interactive robot is only useful if it is capable of displaying a broad range of impedances; indeed, this is one of the central goals set at the outset of this work. Frequently, the robot must be able to represent different impedances in different parts of its workspace (for example a virtual wall features one region that has nominally zero impedance, and a second region that behaves like a stiff spring and/or high-magnitude damper). When viewed in this light, the results of this section are actually quite encouraging. In general, it is found that optimal force control parameters do not vary substantially with the desired stiffness and damping. In other words, a force controller designed for zero impedance should work adequately for active stiffness and damping across a broad range of magnitudes. This is desirable, because the force-error portion of the controller need not make a discrete change in behavior even as the robot end-effector passes from one impedance region to another. The force compensator, which presents the source of the most serious potential stability problem, can remain linear throughout the workspace, even as the motion compensator makes a nonlinear switch between impedance regimes. This may be a benefit of the approach chosen for nonzero impedance design, as opposed to approaches that allow

portions of both the force and motion compensators to vary in designing the best controller to meet a specified impedance. Natural admittance control, for example, concentrates the desired dynamics in the force compensator and requires nonlinear switching of the force controller if the desired impedance changes.

If some parameters are to be fixed across all virtual impedance conditions, but others (such as the DC gain) can be varied, the predicted cost under several conditions can be combined by addition, yielding a composite cost for all the cases together. The set of fixed parameters can then be selected via the lowest composite cost. This is detailed in the next chapter, and an example is provided.

If virtual stiffness and damping are implemented as suggested here, designing a force compensator for the zero-impedance case and subsequently modifying it slightly for the rendering of virtual springs and dampers may be the most effective method of design. Focusing on the zero-impedance case has several advantages. First, this is an important case that is perhaps the most difficult to approximate with high-impedance hardware. Also, because the effect of linear damping appears at low frequencies in this case, the rejection of linear friction is more closely related to the rejection of low-frequency nonlinear friction (like static friction). As shown above, this approximation is less effective when target damping or especially target stiffness is nonzero. This case focuses on reducing the effects of the robot's unwanted friction and inertia, without the added confusion of the structural changes introduced by any of the controller's virtual impedance elements. Friction appears primarily at low to moderate frequencies, and inertia at moderate to high frequencies.

4.3.7 Natural admittance control

The control design algorithm is not restricted to the controller forms specified above; in fact any control law can be incorporated in the model and any subset of its parameters can be searched for the best value. Natural admittance control (described in chapter 2) is an alternative control strategy for interaction, and it is useful to compare its predicted behavior with the results described above.

Although NAC is intended for passive compensation in which the apparent inertia

is no less than the actual inertia, it can be modified to permit the specification of any apparent inertia. For a single-degree-of-freedom example, this is done simply by replacing the actual system mass m with a desired mass M in equation 2.27. This leads to a new expression for the force feedback compensator:

$$G_f = \frac{(m - M)s^2 + (b - G_v - B)s - K}{Ms^2 + Bs + K} \qquad (4.23)$$

m represents the total intrinsic system inertia, where here $m = m_1 + m_2$. $b = b_1 + b_2$ is the total damping to ground. K, B, and M are the target stiffness, damping, and inertia, respectively. G_v is the velocity compensator, which for this example is set to a constant, indicating proportional velocity feedback. Because of the idealizing assumptions in this model, G_v could be stably set at any value if it were applied alone (i.e. without force feedback). This constant was fixed at $G_v = 2000$, a value that performed effectively and stably in robot hardware.

Reducing M below $\frac{m}{2}$ is expected to make the system nonpassive, and introduce coupled stability problems with some environments. Design based on passivity thus excludes such controllers. By instead considering complementary stability, we can predict whether the expected environments are likely to provoke coupled instability with such controlled systems. If not, the apparent inertia can be reduced more aggressively. We can also explore whether the target damping B has any effect on coupled stability.

The search algorithm was extended to search parameters for natural admittance control, with stiffness $K = 0$. The target damping B was varied from 0 to 100 Ns/m (evenly spaced from 0.01 to 100 on a logarithmic scale, with 51 different values including 0), and the target inertia M was varied from 0 to 6 kg (the physical model value) in increments of 0.1 kg. Figure 4-19 shows the lowest value for M that satisfies the complementary stability condition for each value of B. When the desired damping $B = 0$, the coupled system is unstable for all values of M; no point is plotted for zero damping. For all other damping values, the coupled system is stable when $M = 0.5$ kg, but unstable when $M = 0.4$ kg. According to this analysis, damping has no effect.
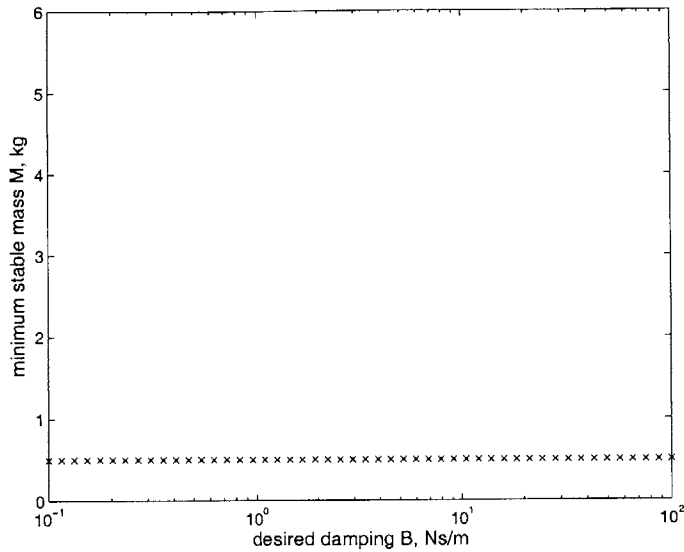
Figure 4-19: Minimum target mass that satisfies complementary stability versus target damping, for natural admittance control.

The best performing controller that satisfies the conditions has the lowest damping, $B = 0.1$ Ns/m, and $M = 0.5$ kg. This outperforms all lag and lead controllers whose design is described in the previous section. The best lead controller that satisfies complementary stability predicts apparent endpoint inertia of 1.1 kg, 120% higher than the 0.5 kg inertia predicted when natural admittance control is used.

If passivity, rather than complementary stability, is used as the stability measure, the stable apparent mass limit varies with target damping. This is shown in figure 4-20, which plots the minimum apparent inertia that makes the system passive for each value of B. When damping is zero, the system is nonpassive for all mass values and is not shown. As damping is increased up to around 6 Ns/m , the inertia can be lowered more aggressively without sacrificing passivity, though never below half the natural value. This dependence on damping is not reported in the early literature on natural admittance control [87, 88], but difficulty has been observed in more recent detailed simulations when the controller seeks to reduce the apparent damping below the level of the physical viscous damping [36]. With virtual damping above 30 Ns/m, the minimum target mass for passivity increases slightly.
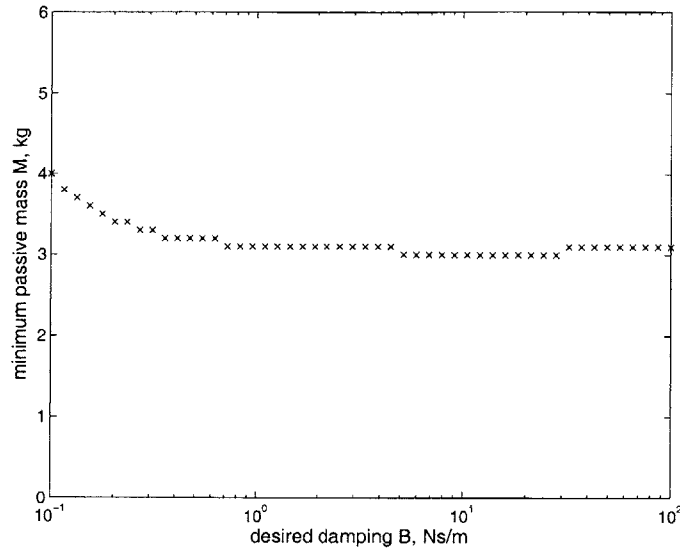
Figure 4-20: Minimum target mass that satisfies passivity versus target damping, with natural admittance control.

The problems with zero damping likely result from the fact that the system attempts to replicate the motion of a marginally stable system in response to a force input. Any error, here probably due to numeric approximation of a continuous system, might render the system nonpassive and a coupled system unstable. This study predicts that natural admittance control is superior to lag-lead control in reducing system inertia. This is interesting though not completely surprising, because natural admittance control uses both motion and force feedback to reduce the impedance, while the previous method is restricted to only force feedback. This study also predicts that passivity is lost for apparent inertia greater than half the physical value if the target damping is not appropriate, as shown in figure 4-20. This is consistent with the analysis of lag and lead control, which suggests that higher target damping can allow higher force gains, but contradicts the analysis in [87], which states that passivity depends entirely on the apparent inertia. On the other hand, the analysis presented here predicts that while apparent damping affects passivity, only the apparent inertia affects complementary stability with the design set of second-order environments. Natural admittance control is presented here primarily for comparison with the newly designed controller, and a detailed analysis of natural admittance con-

173

trol is beyond the scope of this thesis. However, there appear to be some interesting questions surrounding the application of the method to this robot model.

4.4 Implementation

In the previous section, several controllers designed with the complementary stability search method were presented. The method predicts that these controllers can improve performance versus proportional force feedback and reduce inertia more effectively than passive control on a linear robot model. To verify that the method can produce results that are useful on real robotic hardware, several controllers were implemented on the single degree-of-freedom screw-driven linear robot module in figure 2-5. Although some of this robot's physical parameters were used in the design of the controllers, there are a number of differences between the model and the actual robot. Some of these are discussed below in section 4.4.2. Accurately including all of the features of the robot in a model would be almost impossible, and would require a nonlinear model that would be ill-suited for the design tools used in this method. Rather than designing controllers based on a perfect model of the system, the goal here is to see if design with the simple, linearized model presented above can suggest general control strategies and approximate parameters to facilitate improved coupled stability and performance. If useful designs can be obtained from simple models, this may suggest that the method has even greater promise than a method that requires a perfect model, as fully identified systems are seldom available in practice.

The force sensor on the screw-driven module has substantial noise at high frequencies. To attenuate these problems, the force signal passes through a discrete first-order lowpass filter with a break frequency of 30 Hz (188.5 rad/sec). Thus it is not possible to derive any benefit from the lead controller identified above, as its dynamics occur significantly above those of the lowpass filter. A lag compensator was implemented, along with a natural admittance controller.

Lag control was tested with a pole at 0.01 rad/sec, a zero at 6 rad/sec (as suggested by figure 4-10), and several DC gains, including $K_{dc} = 850$, the predicted stability

174

boundary. Also tested were $K_{dc} = 3000$, which is predicted to reject inertia equally as well as a proportional force feedback controller used with this device in physiotherapy [15], an intermediate gain of $K_{dc} = 2000$, and a higher gain of $K_{dc} = 5000$. These higher gains were tested because the algorithm tends to produce conservative gain results (as shown below), and with higher gains the performance can be improved without sacrificing coupled stability. Natural admittance control was implemented with $G_v = 2000$, with target mass varying from $M = 0.35$ to 6 kg and target damping varying from $B = 0$ to 50 Ns/m.

The principal goal of this testing is to compare the stability and performance of the non-passive lag compensator designed with the complementary stability method with a nominally passive natural admittance controller, with the target mass equal to the physical mass (around 6 kg). It is worth reemphasizing that this passive controller represents, to the author's knowledge, the best method available for implementing the currently accepted standard of passivity with confidence in the resulting coupled stability properties. A secondary goal is to compare the test results with predictions from the algorithm for both natural admittance control and lag control, to evaluate the validity of assuming that this model adequately represents the robot. A final goal is to determine which of the proposed controllers has the best stability properties and performance.

4.4.1 Coupled stability testing

To measure the coupled stability properties of the system, interaction with environments dominated by spring behavior was tested. Spring-like environments are usually the most destabilizing for force feedback control [23]. With each controller, the robot handle (end-effector) was placed in contact with each of six compression springs of varying stiffness and a hard plastic block, and stability was determined. Two tests were performed with every combination of controllers and springs. For both tests with all controllers, a virtual spring with a nominal stiffness of 50 N/m was imposed on the robot to provide a tendency for the robot to move into the environment spring. In the first test, the robot handle was brought by hand into stable contact with the
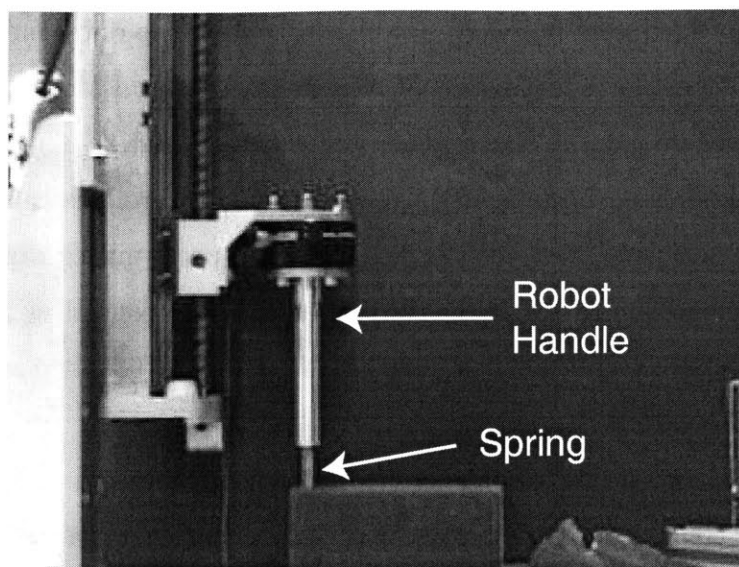
Figure 4-21: Photo of test setup for coupled stability with compression springs.

spring (if possible), and manually given a small force impulse (and then released) to initiate oscillation. If the oscillations decreased in magnitude over time, the system was judged stable for that environment. If they increased, the system was judged unstable. In the second test, the robot handle was released by hand from approximately 3 mm above the point of contact with the spring. If oscillations declined to eventual rest, the system was judged stable for this test, and if they increased or remained constant, the system was judged unstable. This tested the ability of the system to handle transition from free movement to contact. A photograph of the experimental setup is shown in figure 4-21.

The spring constants of the springs, numbered from 1 to 6 in order of increasing stiffness, are $k_1 = 263, k_2 = 976, k_3 = 1782, k_4 = 2652, k_5 = 16,400$, and $k_6 = 57,100$ N/m, respectively. The plastic block has a stiffness of approximately 10^6 N/m (all spring constants were measured in a separate test using the robot's position and force sensors). The mass of the handle assembly outboard of the force transducer was approximately 312 g, including approximately 142 g of outboard mass from the ATI force transducer [34]. Thus the effective environment seen by the robot consists of 312 g inertia, the spring stiffness, plus any damping and inertial properties in the

176

Increasing Environment Stiffness →

	none	1	2	3	4	5	6	block	arm
Kdc=850							XO	++	+
Kdc=2000						X	O	++	+
Kdc=3000					X		O	++	+
Kdc=5000					X		O	+−	−

Controller specified by algorithm   X Maximum stiffness stable in impact test

Explicitly part of design set   O Maximum stiffness stable in contact

Figure 4-22: Coupled stability testing with lag compensators.

spring. This inertia is between the bounds (0.1 to 4.1 kg) specified for controller design. Similarly, the softest spring falls between the stiffness bounds (1 to 401 N/m) specified for controller design. The other environments have higher stiffness and are therefore outside of the design set.

Figures 4-22 through 4-27 show the results of coupled stability testing. There is one figure (figure 4-22) for the four lag compensators, and one figure for each different target mass implemented with natural admittance control. Within each of the figures 4-23 to 4-27, the target damping is varied. Each row in the figures represents a specific controller, described in the left column. Each column represents an environment. The springs are arranged from left to right in order of increasing stiffness. In each row, the "O" symbol marks the environment of highest stiffness that couples stably to the robot, as measured by the first test (initial stable contact), and the "X" symbol marks the highest stable stiffness determined with the second test (impact). If either mark appears on the leftmost column (marked "none"), this means that none of the environments were stable for that test. In the column marked "block," if both tests are unstable there is a dash "–". If both tests are stable, there are two plus signs "++". If the first test is stable but the impact test is unstable, the box is marked with "+-". The final column reports the results of an additional test interacting with a human arm, described in the next section.

All four lag controllers couple stably with the first four springs (up to 2652 N/m) for the impact test, and with all six springs for the "initial stable contact" test. In all the results, the impact test proves more challenging (or equally challenging) than the initial contact test. As the gain is increased, problems arise with impacts with the

	none	1	2	3	4	5	6	block	arm
B=0	XO							--	+
B=0.1	XO							--	+
B=0.3	X		O					--	+
B=0.6		XO						--	+
B=1		XO						--	+
B=2		XO						--	+
B=3		X		O				--	+
B=4			X	O				--	+
B=5			X		O			--	+
B=7				X	O			--	+
B=10					XO			--	+
B=15					XO			--	+
B=20					XO			--	+
B=25					XO			--	+
B=35					X		O	+-	+
B=50						X	O	++	+

Figure 4-23: Coupled stability testing with natural admittance controller, target mass 6 kg.

	none	1	2	3	4	5	6	block	arm
B=0	XO							--	+
B=0.1	XO							--	+
B=0.3	X	O						--	+
B=0.6	X	O						--	+
B=1		XO						--	+
B=2		XO						--	+
B=3		X	O					--	+
B=4			XO					--	+
B=5			XO					--	+
B=7				X	O			--	+
B=10					XO			--	+
B=15					XO			--	+
B=20					XO			--	+
B=25					XO			--	+
B=35						XO		--	+
B=50						XO		--	+

Figure 4-24: Coupled stability testing with natural admittance controller, target mass 2 kg.

	none	1	2	3	4	5	6	block	arm
B=0	XO							--	-
B=0.1	XO							--	-
B=0.3	XO							--	-
B=0.6	XO							--	-
B=1		XO						--	-
B=2		XO						--	-
B=3		XO						--	-
B=4			XO					--	-
B=5			XO					--	-
B=7				XO				--	-
B=10					XO			--	-
B=15					XO			--	-
B=20					XO			--	-
B=25					XO			--	-
B=35						XO		--	-
B=50						XO		--	-

Figure 4-25: Coupled stability testing with natural admittance controller, target mass 1 kg.

	none	1	2	3	4	5	6	block	arm
B=0	XO							--	-
B=0.1	XO							--	-
B=0.3	XO							--	-
B=0.6	X	O						--	-
B=1		XO						--	-
B=2		XO						--	-
B=3			XO					--	-
B=4			XO					--	-
B=5				XO				--	-
B=7					XO			--	-
B=10					XO			--	-
B=15					XO			--	-
B=20					XO			--	-
B=25					XO			--	-
B=35					X	O		--	-
B=50					XO			--	-

Figure 4-26: Coupled stability testing with natural admittance controller, target mass 0.5 kg.

179

	none	1	2	3	4	5	6	block	arm
B=0	XO							--	-
B=0.1	XO							--	-
B=0.3	XO							--	-
B=0.6		XO						--	-
B=1		XO						--	-
B=2		XO						--	-
B=3			XO					--	-
B=4				XO				--	-
B=5				XO				--	-
B=7					XO			--	-
B=10					XO			--	-
B=15					XO			--	-
B=20					XO			--	-
B=25						XO		--	-
B=35						XO		--	-
B=50						XO		--	-

Figure 4-27: Coupled stability testing with natural admittance controller, target mass 0.35 kg.

two stiffest springs. Although no difference between the two highest gains is evident from the spring tests, the highest gain ($K_{dc} = 5000$) encounters problems impacting the block. These results suggest that the block is a less challenging environment than spring 6, and possibly than spring 5. This is plausible, because the block likely has substantially more damping than the coil springs. The effect of environment damping and the apparent conservatism of the algorithm predictions are analyzed in section 4.4.2.

The natural admittance controller testing reveals that the target damping, much more than the target inertia, influences the coupled stability boundary. Increasing target damping helps to stabilize contact that is unstable with lower damping. Surprisingly, substantial reductions in target mass lead to only slight changes in the amount of damping needed to stabilize contact with each spring (on the other hand, higher damping ratio ζ is generally needed to stabilize contact when the mass M is reduced). Indeed, the apparent differences between figures 4-24 through 4-27 may lie within the margins of error of the experiment. This appears to contradict the prediction [87] that only the apparent inertia should influence passivity, and therefore coupled stability. From this test, the apparent damping seems to play a far more important role than the apparent inertia in stabilizing contact with spring-like en-

vironments. The search algorithm predicts some dependence on target damping in determining the passivity of the system (see figure 4-20), but it also predicts a greater dependence on apparent inertia. The properties of this controller appear somewhat different than predicted in the literature, and future study of this may be useful.

Only when $M = 6$ kg and the damping is extremely high ($B = 50$ Ns/m) does the system under natural admittance control interact stably with the block environment (see the last line of figure 4-23). By this test, this is the only controller with coupled stability properties comparable to those of the lag controller with gain as high as $K_{dc} = 3000$.

Interaction with human arms

In addition to the coupled stability testing with spring environments described above, each controller was tested with a human subject using his arm to attempt to destabilize the system. To do this, the subject moved the robot at a variety of speeds and used various arm configurations to provoke instability. As stiffer environments tend to be more problematic than softer, the subject made an effort to stiffen his arm as much as possible, by contracting his muscles, changing posture, and using both arms together to tightly grasp the end-effector. This test is undoubtedly less scientific than the test with springs, as it relies on intentional human behavior that is not perfectly repeatable. Nevertheless, this is the best available representation of the actual environment that the robot will encounter in practice, and therefore should be considered.

When abnormal robot behavior arose, it came in the form of a protracted vibration that tended to worsen with a "stiffer" grasp. An example of the resulting vibration is depicted in figure 4-28. This vibration was the only mode of problematic behavior observed with human arms. A similar vibration emerged with lag control at high gains and with natural admittance control with low target mass. The right most column of figures 4-22 through 4-27 shows whether each controller makes the robot stable or unstable with a human arm. If any vibration like that in figure 4-28 was observed, the controller was ruled unstable for this task and a "-" appears in the right

181
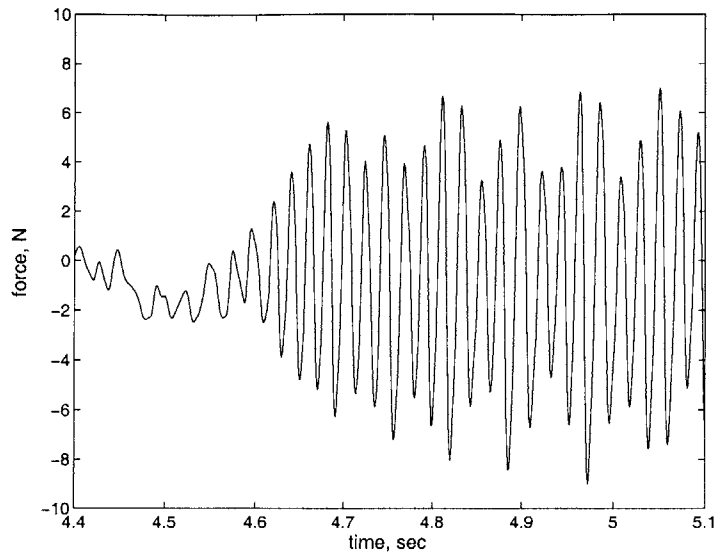
Figure 4-28: Example of vibration that emerges when coupled to human arm for high-gain lag control or natural admittance control with large mass reduction.

column of the appropriate figure. If no vibration was provoked, a "+" is marked in the right column.

For natural admittance control, these results show that avoiding vibration depends on the target mass M; vibration occurs with lower target masses, but not with higher. There is no dependence on the target damping B (at least within the resolution of this test). An observation from the test that does not show up in figures 4-22 through 4-27 is that as the mass is reduced more aggressively with natural admittance control, vibration is induced more easily. With the lowest masses tested, a light grip on the robot handle is sufficient to produce vibration, and if the target mass is made even lower the vibration emerges whenever the end-effector moves. From theory and consideration of the port impedance model, it is not surprising that lowering the apparent mass dramatically is problematic. This also agrees with the output of the controller design algorithm. It *is* surprising that this result does not strongly emerge in testing with spring environments. This could imply that a simple linear spring-mass-damper approximation is not sufficient to capture the relevant behavior of human arms. However, the frequency of the vibration shows that it is far too high

to originate from any voluntary arm activity.

In general, the unstable robot oscillations observed in the tests with spring environments are at a frequency consistent with the spring constant and the apparent mass of the robot. On the other hand, the frequency of vibration when coupled to human arms (35-50 Hz) is far too high to indicate a plausible arm stiffness. It is possible that some unmodeled effect, such as human tissue dynamics (when tightly grasping the handle) or something else is interacting with an unmodeled aspect of robot behavior (e.g. backlash, nonlinear friction, etc.) to produce this vibration.

While the data from coupling to springs may not demonstrate the hazards of reducing mass too aggressively with natural admittance control as clearly as expected, this test does. In either test, the lag compensator effectively maintains coupled stability without unexpected vibration, unless the gain is raised to $K_{dc} = 5000$. When stability with springs, the block, and interaction with arms are considered in aggregate, the tested lag controllers are almost always superior to the tested natural admittance controllers. Only when the desired inertia and damping both increased to their maximum tested values does the natural admittance controller perform as stably as the lag controller with moderate gains ($K_{dc} = 2000$).

Experimental results, whether using springs or a human arm for the environment, show that the controller design algorithm output appears conservative for lag controllers; it appears that higher gains can be attained with the lag compensator in practice than the algorithm predicts (this is analyzed in the next section). This is not true for natural admittance control. This may result from the fact that the latter method depends on both motion and force feedback, as well as an accurate model of intrinsic robot inertia and damping, to reduce apparent impedance, and therefore suffers more from the differences between the robot system and the approximate model used for controller design than does lag control. More detailed analysis on the conservatism of the lag controller follows; a full analysis of natural admittance control is beyond the scope of this work.

4.4.2 Comparison of model and experiment

Because the lag-controlled system couples stably to springs with stiffness much higher than the designed value, and because interaction with a human arm is stable at higher gains than the algorithm predicts, it appears that the results of the model-based algorithm are conservative. Stability when contacting environments outside of the design set, however, is not sufficient to indicate conservatism of the model or algorithm. While not included as a constraint in the controller design, stable interaction with some additional environments is possible and indeed very likely. In guaranteeing stable interaction with all environments in the design set, the controller is likely to also provide stable interaction with a limited set of environments outside the design set.

To determine whether the design algorithm produces conservative results, it is necessary to compare experimental data and model predictions for specific environments. The spring interaction tests provide an opportunity for this. With all robot, control and environment parameters fixed, the coupled system equations can be used to locate the predicted system poles, which determine stability. The environment mass (312 g) was determined by a combination of measurement and estimation. The stiffness is determined by each spring, and was measured with the robot's precision force and motion sensors. The damping of the environment, however, is not as easily determined. But the environment damping, as discussed in section 4.3.5, can be very important for stability. This point is demonstrated in figure 4-29. This plot, determined from the model, shows the maximum stiffness that is predicted to couple stably to the robot under lag control (with $K_{dc} = 3000$) versus the environment damping. The environment mass is fixed at 312 g. If the damping is very low or more than 10 Ns/m, very low environment stiffness should provoke instability. If the damping is at an intermediate level, however, environments with stiffness of more than 3×10^5 N/m can be tolerated. In order to draw a conclusion about the model's prediction of stability or instability with a certain environment, an estimate of the damping must be obtained. This is discussed in the following section.

184
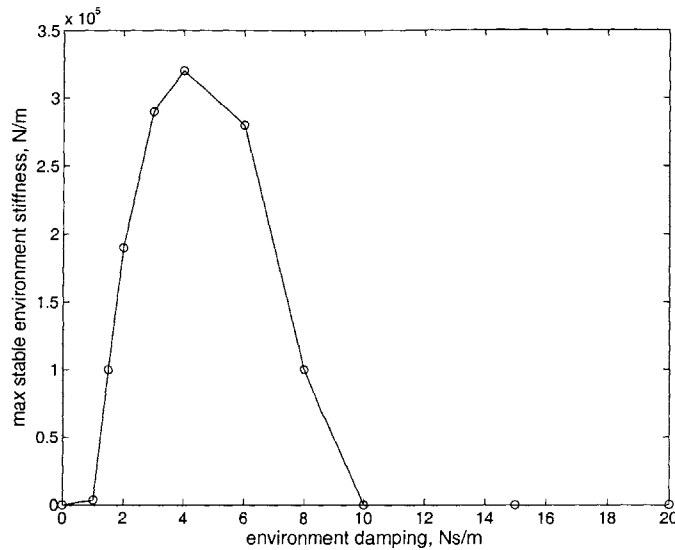
Figure 4-29: Maximum stable environment stiffness versus environment damping, for a lag compensator with $K_{dc} = 3000$.

Figure 4-29 also provides an interesting example of the role of the environment's damping in aiding or compromising coupled stability, as described in section 4.3.5. It appears from figure 4-29 that there is some optimal value of damping that permits the maximum environment stiffness. If damping is less than the optimal value, additional damping can provide a stabilizing influence. If the damping is more than the optimal value, adding more is destabilizing. The role of damping in the environment and throughout the system is critical to stability, but often not obvious.

Estimating environment damping

To estimate the damping in the environment for the coupled stability experiments, the damping rate in each spring must be determined. As the robot handle is much stiffer than any of the springs, almost all of the environment's motion occurs across the spring. Thus the environment's damping is almost entirely due to the structural damping of the spring.

To bound an estimate of the damping coefficient for each compression spring, a simple test was designed. The setup for this test is sketched in figure 4-30. A weight
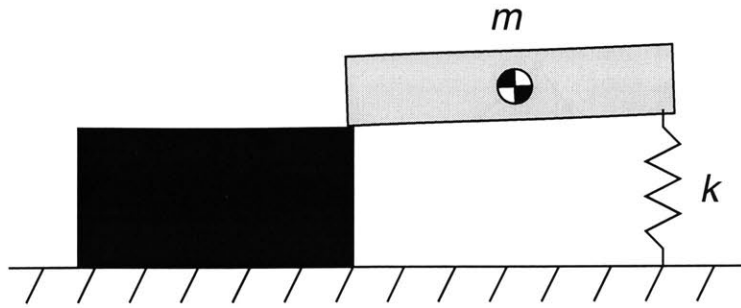
Figure 4-30: Schematic of test setup to determine damping ratio of compression springs.

with known mass m was placed with one end supported on a block and the other by the spring. This creates a simple resonator between the mass and the spring's compression mode. Because the center of gravity of the mass moved vertically about half as much as the end of the spring, and because displacements were small in this test, the effective inertia for this system was assumed to be $\frac{m}{2}$ (If rotational inertia is included, this estimate is slightly larger; however the effect on the final results is minimal compared to the assumed error of estimate). The oscillation of the resonator is governed by this inertia, the known spring stiffness k, and the unknown spring damping. Different masses were used for the different springs to keep the resonant frequency low enough for oscillation to be easily observed.

Oscillation was initiated by striking the mass with a small hammer. Two visual observations were made for each test. The first was a measurement of the amount of time that visible oscillations continued. The second was an estimate of the approximate amount of time for the amplitude of oscillation to decay to half of its original size. These two observations were used to obtain four estimates of the damping ratio ζ. Assuming that the magnitude of the oscillation must exceed either 1%, 2%, or 5% of the initial magnitude to remain visible, three estimates of ζ were obtained from the first measurement using a damped linear model. A fourth value for ζ was obtained from the estimate of the time to decay to half the oscillation amplitude. From the estimates of ζ and the known stiffness and inertia, the estimated damping coefficient

b can be computed with the following equation:

$$b = 2\zeta\sqrt{k\frac{m}{2}} \tag{4.24}$$

This method provided measurable data for the four softest springs; the two stiffest springs displaced too little to visually observe significant oscillation. All four springs are quite lightly damped; each tended to oscillate visually for between 5 and 30 seconds, at frequencies greater than 5 Hz. As a result, all estimates of ζ and b are very low, and fairly similar given the crude test method. The mean b_m of the four estimates of b for each spring was computed. An upper bound b_{mu} was computed from the equation:

$$b_{mu} = b_m\sqrt{10} \tag{4.25}$$

and a lower bound b_{ml} from the equation:

$$b_{ml} = \frac{b_m}{\sqrt{10}} \tag{4.26}$$

Thus the final estimate of b for each spring can vary by an order of magnitude between b_{ml} and b_{mu}. Figure 4-31 plots the range of damping estimates for each spring. The error bar indicates the region from b_{ml} to b_{mu}. The data points show the four individual estimates of b for each spring. In each case the bounded estimate includes a wider set than all four individual estimates.

Comparing coupled stability results to predictions

With upper and lower estimates on the damping coefficient for each of the four softest springs, the model stability predictions can be compared to the experimental data. The model was used to predict coupled stability with each spring, based on both the lower and upper bound of the damping coefficient. If the model predicted that either damping value would stabilize interaction, the coupled system was judged stable; this ensures that any possibility of the model predicting stability is recognized.

Figure 4-32 is a reproduction of figure 4-22 that includes the model predictions
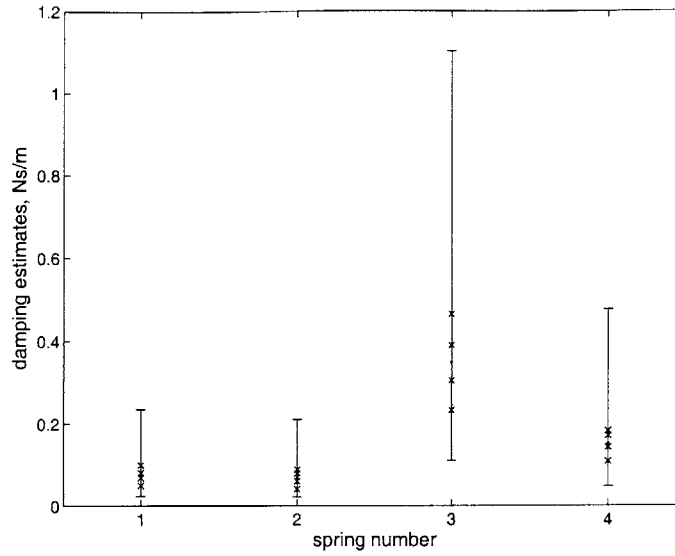
Figure 4-31: Damping coefficient estimates for each of first four springs. Individual estimates are shown with an "x". The mean estimate of each spring appears as a dot. The error bars indicate the upper and lower bounds used for coupled stability predictions.

of the maximum spring stiffness that contacts the system stably, marked with a "P". When the gain is $K_{dc} = 850$, the model predicts stable coupling to all six springs (this is true if the damping is any non-negative value less than 70 Ns/m, including zero). However, it predicts unstable coupling to the block, regardless of damping. The experiments show stable coupling to all springs as well as the block. For each of the higher values of K_{dc} tested, the model predicts instability even with the softest spring, while the experiment shows that each is stable up to at least the fourth spring. The model produces conservative results; that is, many cases are identified where the model predicts instability but the experiment shows stability.

Differences between robot and model

The previous section shows that the design algorithm and model produce output that is conservative (though less conservative than passive controllers). In fact, this is evident even independent of any novel control design. From figure 4-9, placing a compensator with a pole around 200 rad/sec in the force feedback path of the

188

Increasing Environment Stiffness →

	none	1	2	3	4	5	6	block	arm
Kdc=850							PXO	++	+
Kdc=2000	P					X	O	++	+
Kdc=3000	P				X		O	++	+
Kdc=5000	P				X		O	+−	−

Controller specified by algorithm

▨ Explicitly part of design set

X Maximum stiffness stable in impact test
O Maximum stiffness stable in contact
P Predicted maximum stable stiffness
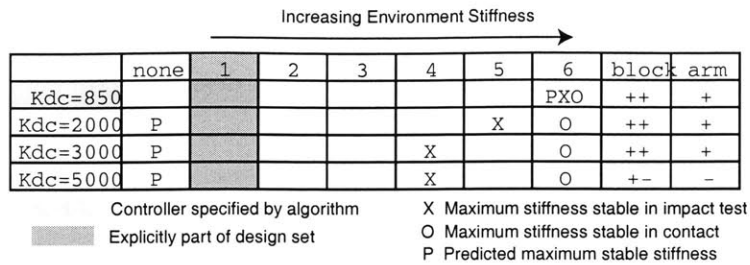
Figure 4-32: Predicted and measured coupled stability boundaries, lag control.

model is predicted to almost preclude the reduction of apparent impedance via force feedback. However, on the actual robot, the filter on the force signal contributes such a pole. Nonetheless, force feedback is implemented with modest gains without the anticipated coupled stability problems [15, 16]. This suggests that there is some feature lacking in the model but present in the robot that makes the physical system more amenable to force feedback control. One unmodeled characteristic that might perform this function is the large amount of Coulomb friction. This could act to "drain" some of the excess energy that is generated by non-collocated force feedback control and that tends to destabilize contact.

There are a number of other effects that are known to exist in the robot but are unmodeled. Control is implemented discretely with some (small) computation delay, but modeled as ideal continuous control. The system has moderate amounts of backlash in the screw. The actuator and servoamplifier have bandwidth limitations, and there are countless unmodeled structural modes of vibration. The sensors have limitations as well. Some combination of these may lead to the vibration (around 35-50 Hz) observed in the robot under various forms of high-gain force feedback control when contacting human limbs, as documented in the coupled stability tests described above. This vibration is not predicted by the model.

In spite of these simplifications, the linear model is adequate in that it facilitates design of a controller that provides improved stability properties, as shown in the section above, and improved performance, as shown in the next section. Although the controller is conservative, it is less conservative than controllers designed for

189

passivity. It improves the state of the art by reducing conservativism, but remains sufficiently conservative to guarantee stability when desired.

4.4.3 Performance testing

Evaluating performance, as defined throughout this thesis, of an interactive robot means evaluating its impedance. This can pose a significant challenge, as the system may go unstable when certain means of testing are applied. For instance, a common approach to testing impedance is to impose motion with specified frequency content and measure the force response. Because the controllers tested here are not designed to interact stably when interacting with a pure motion source, this is likely to introduce instability (indeed, this is predicted in many cases; see appendix A). Therefore other means must be devised to evaluate performance. Several approaches and their results are described below.

Perhaps the best way to evaluate performance, or "feel," is for a human to physically interact with the robot. Humans are skilled at identifying what feels easy to move, if not at quantifying the same. The feel of the robot has substantially assisted the evaluation of performance, and comments on this subject are added to interpretation of the quantitative data where appropriate.

Stiffness, Coulomb friction

In order to evaluate the low-frequency friction observed in real interaction with a human arm, a virtual spring was imposed with each of several controllers. With the virtual spring in place, the system was actuated very slowly by hand, first up and then down. Care was taken to keep velocity and acceleration very low; the velocity never exceeded 0.05 m/sec. A similar test provided helpful information in [15].

This test was performed with a simple impedance controller (a PD controller without force feedback), a (theoretically passive) natural admittance controller with target mass equal to the actual system mass (6 kg), and the lag compensator designed by the search algorithm. The resulting relationship between force and position for
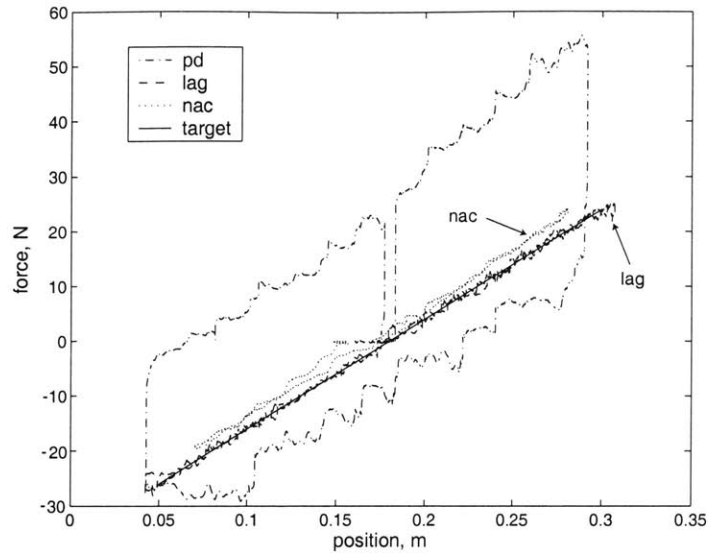
Figure 4-33: Force versus position of the robot representing a virtual spring. PD control, lag force feedback control, and natural admittance control shown along with target behavior. Robot is actuated by hand at speeds below 0.05 m/s.

each controller, along with the target stiffness of 200 N/m, is shown in figure 4-33. If the system has no friction at zero or extremely low velocities (i.e. no static or Coulomb friction), and if the stiffness is accurately implemented, the trace should follow almost exactly the target stiffness curve. The effects of inertia and viscous damping are negligible, because of the low velocities and accelerations. The simple impedance controller does nothing to compensate for the Coulomb friction in the system. It is evident that the system has as much as 20 N of Coulomb friction. The friction is also periodic with the screw's rotation, as is evident from the periodic variation in the curve (the screw lead is just under 0.02 m). The center of this curve is shifted up from the target curve because this controller does not use gravity compensation, and it takes more force (approximately 10 N) to move the system up than down.

The lag compensator rejects almost all of the low frequency friction. Of the three controllers shown, this curve best tracks the target stiffness. The only deficiency in its performance in this task can be seen by blowing up a portion of figure 4-33, shown in figure 4-34. There is noticeable high-frequency content in the lag compensator

191
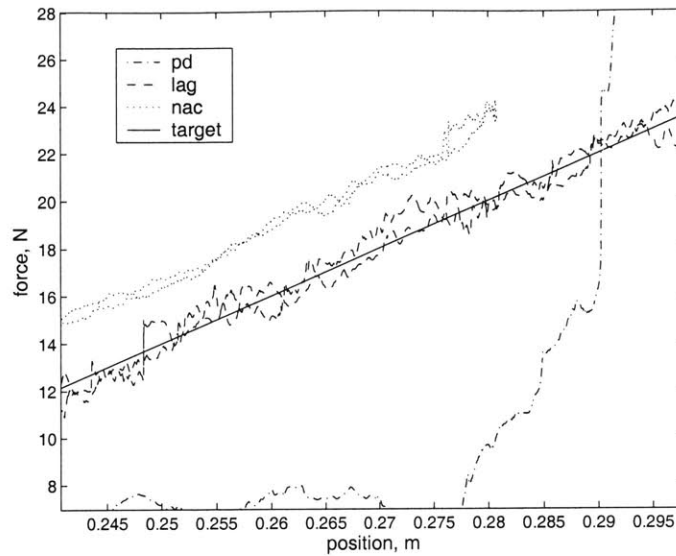
Figure 4-34: Closeup of figure 4-33.

trace. This is consistent with the feel of the system, which is slightly "jerky" at very low speeds. This is likely due to the interaction of the low-frequency lag action in the controller (that approximates integral action) with Coulomb friction, which is discussed below. Still, friction rejection is quite effective, and qualitatively, the uneven feel is only a minor nuisance.

The natural admittance controller effectively rejects most of the static and Coulomb friction. This controller suffers from several other problems due to the fact that its motion feedback is based on velocity, not position. The controller directs the system to move with a specified velocity trajectory in response to the applied force. Errors in tracking this trajectory are integrated to position, so the virtual trajectory, in this case a stationary point, tends to drift as time goes on and errors collect. This is why the curve does not track the straight target line as effectively as the lag compensator. This controller also has difficulty rendering the desired stiffness accurately; in fact, this test required a nominal stiffness of 170 N/m in order to achieve an actual stiffness close to 200 N/m. The natural admittance controller is at its worst in static or near-static situations, where the velocity is very low (and where velocity estimation from position data is poor). These problems are not insurmountable, but require

192

additional measures to correct (for example, position feedback that occasionally corrects accumulated errors). On the other hand, the curve is considerably smoother than the lag compensator, as can be seen in the zoomed in plot in figure 4-34. This is consistent with the feel of the system; the natural admittance controller feels like it rejects friction better than the lag compensator, and it feels more like a natural physical system.

Inertia

A second test was performed to estimate the apparent inertia of the system under control. A virtual spring was again imposed by each control algorithm. A mass with magnitude 1.55 kg was then hung from the end-effector with monofilament line. The force (15.2 N) applied by the mass caused a displacement against the virtual spring. The string holding the mass was then cut, freeing the end-effector to rebound to its reference position. With sufficient friction rejection, the system is underdamped and the handle oscillates about the reference position. This is a method for measuring the system's initial condition response. Because the stiffness of the virtual spring is known, the frequency of oscillation about the reference position reveals the apparent endpoint inertia.

With natural admittance control, a specific controller objective is to replicate the behavior of a second-order linear system, and high-gain force and motion feedback are used to this end. The result is a system that behaves very much like a linear system. An example of the initial condition response with a natural admittance controller is shown in figure 4-35. The controller was used to impose a nominal stiffness of 170 N/m (shown in the previous experiment to approach 200 N/m statically), with target inertia of 6 kg and target damping of 10 Ns/m. The initial condition response of a linear second-order system with inertia of 6.1 kg, damping of 10.5 Ns/m, and stiffness of 200 N/m is plotted with the test data. The inertia and damping parameters were determined by a visual fit to the response. Although the static stiffness is inaccurate, the controller closely replicates the target inertia and damping, and the response is as predicted. This test procedure was repeated for several different target mass values.
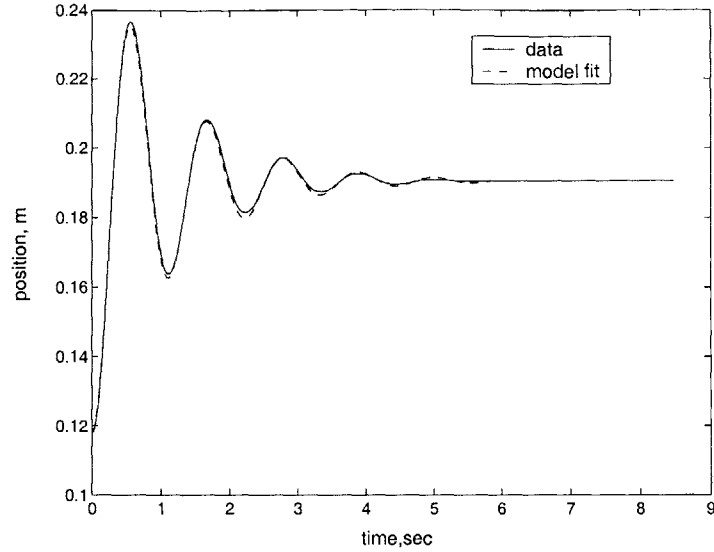
Figure 4-35: Sample initial condition response data and fit with natural admittance control. Target inertia is 6 kg.

Target Inertia (kg)	Measured Inertia (kg)	Percent Error %
6	6.1	2
2	2.2	10
1	1.15	15
0.5	0.66	32

Table 4.2: Results of initial condition response tests to determine apparent robot inertia with natural admittance control.

Table 4.2 provides the results of this test. The measured mass is consistently larger than the target mass. As the target mass decreases, the error increases. As mentioned above, the dynamic performance of natural admittance control is much better than its static performance. It shows a tendency to drift with the mass hanging, so that after oscillation it settles to a location that differs from the original center of its spring. While the dynamic trajectory in figure 4-35 fits the model very accurately, the controller's *static* performance remains problematic.

Lag control boosts the low-frequency gain to reduce low-frequency friction. Like the natural admittance controller, the objective is to suppress unwanted friction and inertia, but this approach does not explicitly make the system track trajectories

194

typical of a linear system. This is an important difference. Using motion control, the natural admittance controller overrides the undesired system behavior and makes the system tend toward a specified linear response. In contrast, the ideal behavior of the lag compensator is that of a pure spring (in this case, as the target damping is zero). Once the input force is removed, the end-effector position should return immediately to the center position of the spring. In the time domain, any deviations from this perfect step trajectory are treated as errors. The low-frequency lag control resembles proportional-integral action, and interacts with the system's nonlinear friction to bring the trajectory close to the step.

Figure 4-36 shows an example of the initial condition response of the system under lag control with a DC gain of $K_{dc} = 2000$, and the other control parameters as above. In comparison to natural admittance control, this response is significantly nonlinear. The system does not oscillate for long, despite the fact that friction is rejected effectively (as shown in the preceding section). A linear model was visually fit to the system in figure 4-36 with the primary goal of fitting the initial rise and overshoot, which is sufficient to generate an estimate of inertia with known stiffness. The fit displayed in the figure uses inertia of 1.75 kg and linear damping of 7 Ns/m. This inertia is 30% larger than the value that the robot model predicts for this controller (1.35 kg). Still, it is a 42% reduction below the theoretical passive limit of 3 kg. The procedure was repeated for initial condition responses with several different gains. Table 4.3 shows the predicted and measured inertia for each, and the percent error. The measured inertia is consistently larger than the estimated inertia. The errors are slightly larger than they are for natural admittance control, as shown in table 4.2. This could result from several factors, including errors in the system model parameters (particularly the total inertia or natural frequency), or errors from fitting a linear model to a nonlinear response. Still, inertia is rejected significantly below the 3 kg theoretical passive limit with each of these controllers.

The lag controller experienced no stability problems in the course of testing. With the natural admittance controller, when the target mass was reduced to the lower levels (i.e. 1 kg or less), the system sometimes initiated oscillations when the mass
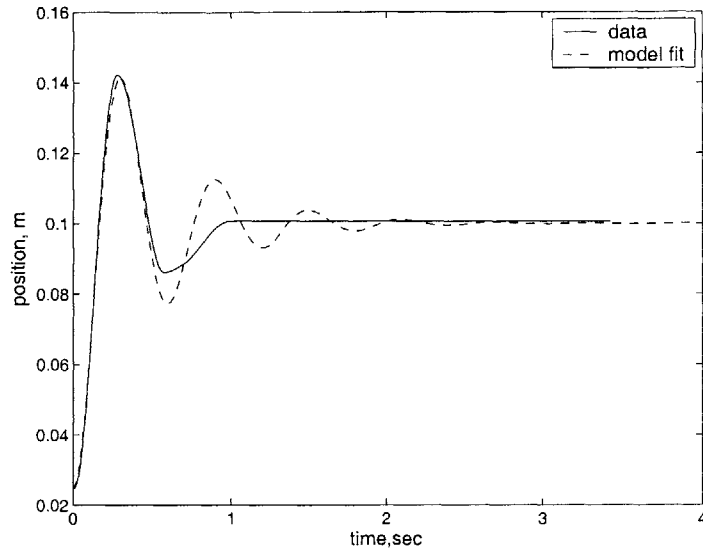
Figure 4-36: Results of initial condition response tests of apparent robot inertia with lag control.

K_{dc}	Predicted Inertia (kg)	Measured Inertia (kg)	Percent Error %
2000	1.35	1.75	30
3000	0.95	1.2	26
5000	0.6	0.95	58

Table 4.3: Results of initial condition response tests to determine apparent robot inertia with lag control.

was hung from it, after being settled to rest. This may indicate that a mass connected to the robot through a monofilament line is another problematic environment for natural admittance control.

4.4.4 Conclusions from initial results

The simulations presented in the first part of this chapter show that the search algorithm based on the proposed complementary stability criterion and a port function measure of performance can converge to a non-obvious solution that improves performance versus state of the art passive controllers when applied to a single-resonance robot model motivated by the literature. The design method leverages partial knowledge of the environment to limit the destabilizing effects of force feedback where necessary and to reduce impedance more aggressively where the environment permits it. Given a common yet flexible control structure, the algorithm can distinguish the benefits of various force control strategies as proposed in the literature, and select and parametrize the best form for the modeled problem.

The second part of this chapter shows that a controller designed with a linear model of the robot and environment can be applied to a more complicated nonlinear system and can improve stability and performance for interaction. The algorithm provides a nonpassive lag controller that provides coupled stability properties that are superior to natural admittance control and that effectively reduces both friction and inertia. Friction rejection is dramatically superior to proportional control alone, and the improvement is due to the controller form produced by the design algorithm. This improvement is perhaps best realized by comparing figure 4-33 with figure 4-37 (reproduced from figure 2-25 for convenience), both of which show force feedback on the same robotic hardware. The lag controller accurately implements the desired stiffness with almost no low-frequency friction. The algorithm also provides the critical break frequency of the zero in the lag compensator, to trade off friction and higher-frequency inertia rejection. The model differs from the robot in several ways, but it suffices to facilitate development of controllers that improve stability and performance.
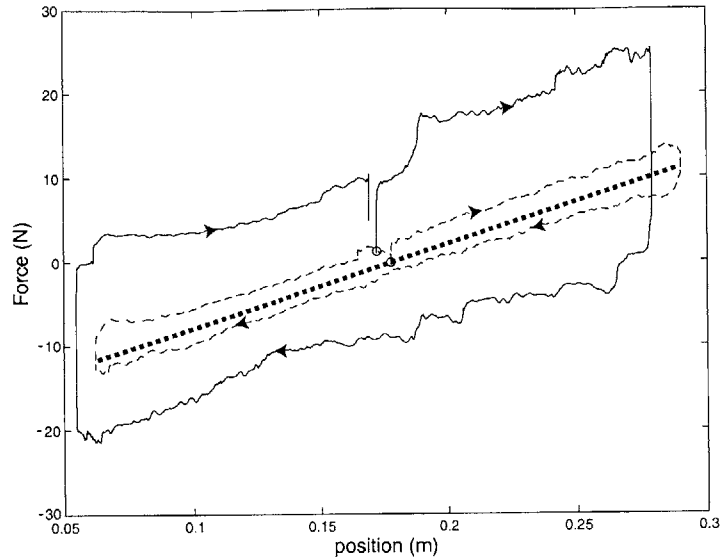
Figure 4-37: Reproduction of figure 2-25. Virtual 200 N/m stiffness represented on the screw-driven robot from example 2.2 with simple impedance control (solid line) and proportional force feedback control with $K_f = 5$ (light dashed line). Actuated by hand. Ideal behavior shown by heavy dashed line.

In spite of these improvements, force feedback controllers still pose a stability risk when they are used to reduce apparent impedance. In the example, apparent inertia can be reduced significantly more than the 50% limit presented by passive control, but the reduction is still less than an order of magnitude (for example, the lag controller with $K_{dc} = 3000$ reduces the inertia by around 80%). Systems that require actuators with very high force densities may need high-ratio gear reductions to meet the force and weight requirements, and the resulting inertia may be several orders of magnitude too high. If control does not suffice, the only remaining option is to alter the physical system structure. The next chapter applies the tools developed in this and the previous chapters to this task.

198

Chapter 5

Series Dynamics to Improve Interaction

As shown in the previous two chapters, force feedback can reduce apparent impedance; by taking advantage of limited knowledge of the environment's dynamics, undesired impedance can be rejected more effectively without jeopardizing coupled stability. Still, the capabilities of this approach are fundamentally limited. As the examples show, if the environment has predictably bounded impedance properties inertia can be reduced by more than the fifty percent limit required to preserve passivity, but the apparent inertia usually remains within an order of magnitude of the physical value. This improvement may be enough to solve the high-force density, low-impedance actuation problem for some applications; a moderate gear ratio could be used to reduce actuator mass without increasing the reflected impedance so much that moderate force feedback, sculpted to take advantage of known dynamic properties of the environment, cannot lower the apparent impedance to the specified level. Hardware must be intrinsically backdrivable with intrinsic impedance that is, as a rough rule of thumb, at least within an order of magnitude of the desired value.

This approach is not likely to satisfy all needs. The impedance of the screw robot module, starting with around 20 N of undesired Coulomb friction and 6 kg endpoint inertia, can be brought to the desired impedance level with force feedback, as shown in the previous chapter. However, it does not exceed the target performance by

much, particularly with respect to inertia, regardless of the specific control implementation. This is the case in spite of the fact that the machine uses an intrinsically low-impedance screw transmission and only a modest gear reduction (in the form of a relatively long screw lead)[15]. Because the gear reduction is limited, the actuator mass is substantial, and including the mass of the transmission structure, the module mass approaches 8 kg. This is four times the target mass of 2 kg for this module (intended for mounting on the links of a low-impedance planar robot). Increasing the gear reduction to reduce the mass would likely drive the impedance too high to be brought within specification by force feedback. As mentioned in chapter 2, in some applications the real actuator weight may be even more critical, and those applications may not include any means by which to compensate for weight.

The apparent limits of force feedback leave us searching for an alternative way to make the apparent impedance of a high-impedance device closer to its target. Fortunately, another method is available to partially suppress the appearance of the unwanted high-impedance dynamics at the port of interaction. This strategy involves making changes to the physical robot system to make it appear more like the desired impedance, which can be done by placing additional dynamics in series between a high-impedance actuator and the environment. Depending on the details of the robot and series dynamics, these dynamics can either dominate the apparent dynamics presented to the environment, or simply transmit slightly modified (or filtered) robot dynamics.

5.1   Benefits of series dynamics

"Soft" dynamics are commonly placed in series between a robot that uses force feedback and an environment to help stabilize the interface. Variants of this approach have been used since the earliest attempts at robot force control [118], for example in the form of compliant pads on the robot end-effector. This approach makes sense considering that such robots generally have stability problems with extremely stiff environments; the pads make everything in the environment appear softer to the robot.

Similarly, Eppinger [38] recommended using a remote center compliance (RCC) device [37] for this purpose. In each case, a passive device is used to provide compliance near the endpoint of a robot to help stabilize interaction.

Soft dynamics have their cost; the placement of compliance outside the control loop can reduce both bandwidth and the ability to closely regulate position. Conversely, benefits can also be derived from series dynamics that are not springlike; below it is shown that inertial dynamics can offer advantages. The benefits of series dynamics come from several sources, and it is important to understand each separately before determining the best dynamic structure to use for an application.

5.1.1 Changing physical system structure

One way that series dynamics can help reduce apparent impedance is by altering the physical dynamics of the system to more closely resemble that of the desired system. The examples in the previous chapter demonstrate the difficulty of attempting to make a system that most closely resembles a damped inertia behave like a damped spring or a zero-impedance force source. If the system is always expected to behave like a spring, it may be possible to meet this requirement by endowing the system with the physical properties of a spring. This point can be made by looking at the port impedance of two simple systems, one that uses soft series dynamics and one that does not. If the friction in figure 2-24 is linear damping to ground, then:

$$F_f(x, \dot{x}) = b\dot{x} \qquad (5.1)$$

A simple impedance controller with stiffness K and damping rate B can be applied to this system with the control law:

$$F_a = -Kx - B\dot{x} \qquad (5.2)$$

201

The port impedance of this system is:

$$Z(s) = \frac{F_e}{\dot{x}}(s) = ms + (b + B) + \frac{k}{s} \tag{5.3}$$

If a spring with stiffness k_s and damper with damping rate b_s are added to this system between the mass and the environment, the resulting system is pictured in figure 5-1. A simple impedance controller with the following stiffness (this stiffness is selected such that the total series stiffness is equal to K at low frequencies):

$$K_2 = \frac{K k_s}{k_s - K} \tag{5.4}$$

and damping rate B is applied with the following control law:

$$F_a = -K_2 x - B\dot{x} \tag{5.5}$$

The port impedance magnitudes $|\frac{F_e}{x}|$ for the first system and $|\frac{F_e}{x_2}|$ for the second system are plotted in figure 5-2. The parameter values are $m = 6$ kg, $b = 25$ Ns/m, $k_s = 250$ N/m and $b_s = 1$ Ns/m for the physical system, and $K = 200$ N/m and $B = 1$ Ns/m for the target and control parameters. The figure shows the target impedance along with both curves. Although both systems accurately provide stiffness at low frequencies, the system without the series elements has strong inertial properties that dominate above 7 rad/sec (the damping $(b + B)$ determines the damping ratio of this break point). At high frequencies, the system with series dynamics behaves much more like the target system, as the damping dominates. At all frequencies, the series dynamics dominate. This is because the stiffness of the series spring, 250 N/m, is close to the target value of 200 N/m, so the controller spring is very stiff ($K_2 = 1000$ N/m), and the movement amplitude of the high-impedance mass x_1 is much less than the movement at the endpoint x_2. Because the mass does not move as much, the associated damping and inertia play a much smaller role in the endpoint dynamics. Because the dynamics closest to the endpoint are close to the desired behavior, the control system's role in presenting the correct impedance (in spite of
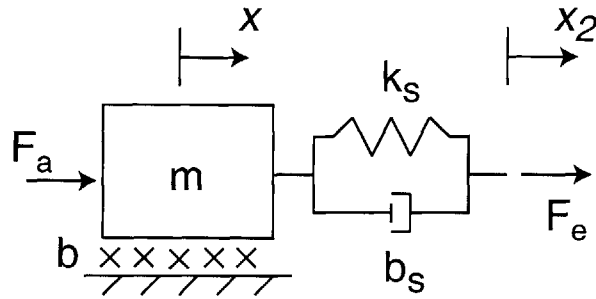
Figure 5-1: Simple damped-mass robot model with a parallel spring-damper in series between robot and environment.

the high-impedance mass) is minimized.

Another way of thinking about the same point is to envision "decoupling" the motion of the high-impedance source from the motion of the endpoint. If the source inertia is retarded by static friction, for example, the addition of a compliant member between this friction and the environment permits the environment to move, at least with small forces and displacement, without breaking the static friction. In this arrangement, the troublesome nonlinear friction is no longer exactly tied to the motion of the environment. Similarly, the inertia m is decoupled from the endpoint motion x_2. In the system in figure 5-1, beyond the break frequency of the series spring and damper the motion x_2 is almost completely decoupled from x_1, and therefore the inertia is insignificant at the endpoint. This point is reiterated below in the context of discussing the optimal characteristics for series dynamics.

This first benefit of soft series dynamics for human-interactive applications is purely mechanical in nature; it holds true with or without force feedback. Assuming that for human interaction, spring-like robot behavior is preferred over inertial and friction behavior (this is generally used in practice; see [56]), it is advantageous to make robot port behavior intrinsically more spring-like, and this can be accomplished by adding the described dynamics. Although the point is somewhat elementary, it is worth stating that force feedback is not necessary to derive benefit from these purely physical changes.
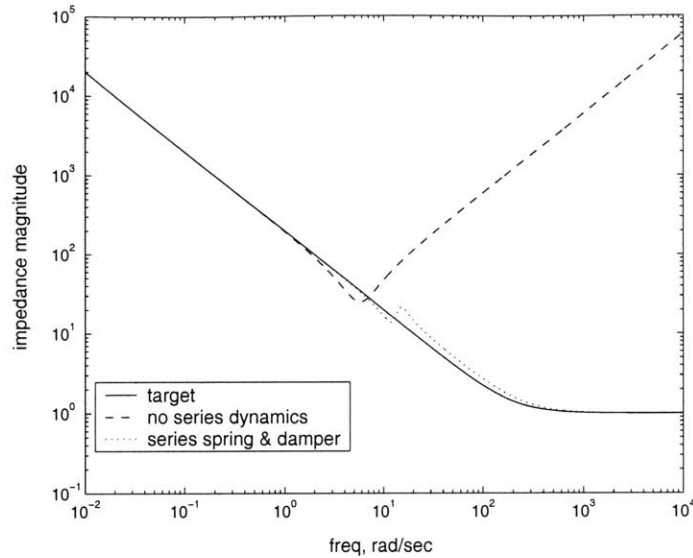
Figure 5-2: Port impedance magnitude of robot in figure 5-1, with and without the spring-damper in series. Target behavior is also plotted.

5.1.2 Stabilizing force-feedback systems

The second important way that series dynamics can help bring port impedance closer to the target impedance is by stabilizing contact between a robot using force-based control and its environment that would otherwise be unstable. This can permit more aggressive force feedback control to reduce unwanted impedance properties.

From the coupled stability experiments in the preceding chapter, it is clear that stiff spring-like environments tend to provoke instability when coupled to robots that employ force feedback to reduce apparent impedance, with stiffer springs causing more problems than soft springs. If a spring of stiffness k is mounted to the robot such that it becomes the new point of contact with the environment, this appears to the robot no different than if each environment has the properties of this spring in series with its intrinsic properties. If $k_1 > k_2$, and contact with k_1 destabilizes a robot while contact with k_2 does not, k can be selected such that $k < \frac{k_1 k_2}{k_1 - k_2}$. Hence contact by the new robot system (including the spring of stiffness k) with an environment of stiffness k_1 appears to the robot as contact with a spring of stiffness $< k_2$, which is stable. The exact same actuation, sensing, and control law is used for both cases, but the

204

first is unstable when contacting k_1, while the second is stable. This is accomplished simply by physically adding a spring to the robot.

Below it is shown that if properly placed damping is included in the series dynamics, this can provide a versatile means of stabilizing contact with a broad range of environments by passivating a nonpassive port [35]. The stabilizing influence of soft series dynamics has been studied and exploited by several research groups, and is applied systematically below using the design methods developed in the previous two chapters. Alternative design suggestions from the literature are first provided for context. The search algorithm tools are then applied to design systems similar to those suggested in the literature in a systematic way. This reveals interesting new results that help to explain the successes and limitations of prior work. In addition, inertial series dynamics are explored and found to have benefits as well.

5.1.3   Design strategies for series dynamics

Several pieces from the recent literature suggest methods for incorporating soft series dynamics in actuators for interaction. The state-of-the-art comes from the related fields of force-controlled actuation and passive robot control, and an approach from each is described below.

Series elastic actuators

Series elastic actuators [95, 99, 100] use a spring mounted in series to improve performance of a force-controlled actuator. Designed for use with highly-geared source actuators, series elasticity protects the gears from high impact loads and permits the storage of energy in the springs. The spring can also improve closed-loop force control performance by trading (possibly unnecessary) high-force bandwidth for improved low-impedance performance. The analysis provided by the authors uses a lumped mass for the actuator and geartrain, and therefore does not demonstrate the non-collocation problems caused by feeding back the endpoint force. Early analysis [95] depended on the assumptions that the geared actuator and environment both
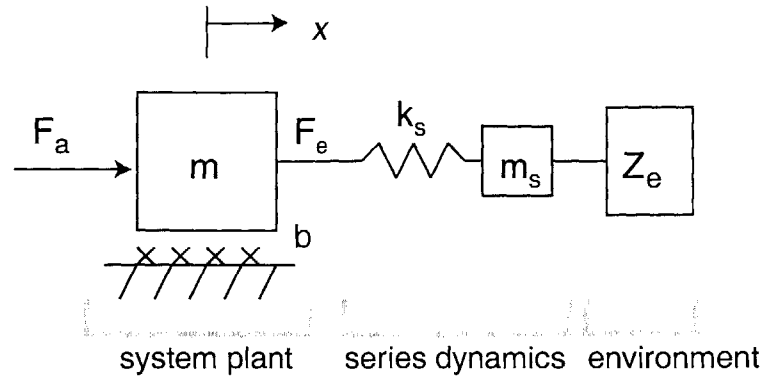
Figure 5-3: Simple robot model with a series elastic element and a series mass, connected to an environment.

have high impedance relative to the spring element so that the problem could be treated as position control of spring deflection, rather than force control. This work suggested PI force control, where the integral action causes passivity to be lost, and coupled stability to be compromised. The authors recognize, however, that environments with high natural frequencies are problematic for force control. They suggest a minimum mass at the output of the spring (generally produced automatically from the actuator structure) to bring the interface natural frequency down to a frequency below the force controller bandwidth. The rationale behind this is that within its bandwidth, the force controller can make the actuator appear to have zero stiffness, and therefore can prevent any oscillation that might create an unstable mode. Appending series dynamics like this to a linear, single-mass actuator model results in the structure shown in figure 5-3, where the spring k_s is the critical element and the mass m_s is due to the robot structure.

More recent analysis [99, 100] considers the inertial and friction dynamics of the drivetrain along with the dynamics of the series elastic element, and recognizes the desire to minimize actuator and gear stiffness. This modeling approach is more consistent with the idea of creating backdrivable hardware and augmenting performance with control. PD force control is suggested in order to regulate the closed-loop natural frequency and damping ratio.

To select k_s for a series elastic actuator, Robinson [100] suggests first sizing the

motor and geartrain on the basis of force, speed, and power requirements. A target high-force bandwidth is then selected, and this places a lower bound on the spring constant independent of control. Design levels for static friction and for the impedance due to the drivetrain mass and damping are selected, and because decreasing k_s decreases both of these quantities, an upper bound for the spring constant is determined. A value between these two bounds is selected, and a PD force controller is then designed to satisfy the bandwidth and damping ratio specifications (of the force transmission transfer function). Of course, the higher the control gains can be raised, the better friction and mass can be rejected. If suitably high gains cannot be stably achieved, it may be necessary to select a softer spring (at the expense of bandwidth) and redesign the gains. No guidelines are provided for the sizing of the mass m_s, which is not mentioned in the more recent work.

Unfortunately the published design procedure does not provide a systematic way to predict maximum force gains that are likely to permit stable interaction (through the spring) with the environment, or to estimate the amount of impedance rejection that can be tolerated. This deficiency results mainly from the fact that a collocated model is used that cannot predict the coupled stability loss due to force feedback. As a result, series elastic actuator design may likely include design iterations in physical hardware using different spring constants. Nonetheless, the method correctly identifies that the presence of a series elastic element can both reduce intrinsic system impedance (at the cost of bandwidth) and stabilize force-controlled systems.

Mechanical filters for passivity

An alternative architecture for "mechanical filters" uses a spring and damper, in parallel with each other but together in series between the robot end-effector and the environment [35]. This is shown in figure 5-4. This model is similar to that in figure 4-1. The difference is that the force transducer spring and damper, k_f and b_f, have been replaced with the series mechanical filter spring and damper, k_s and b_s, which are design parameters rather than fixed values. Because the force transducer stiffness is much greater than the filter stiffness ($k_f \gg k_s$), for convenience the force
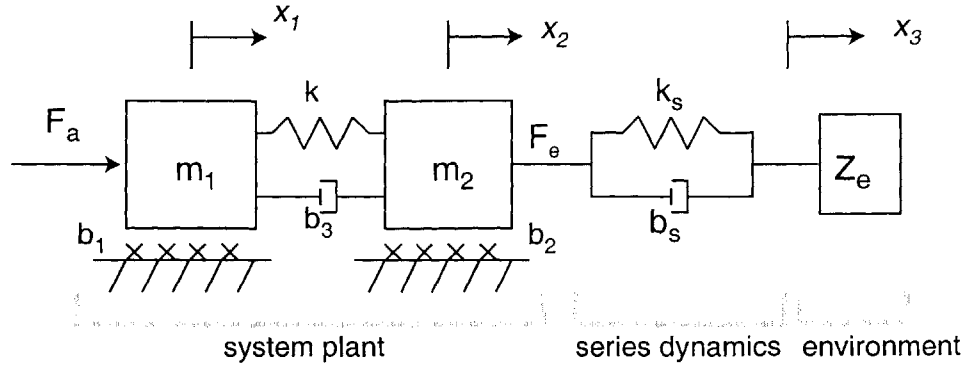
Figure 5-4: Single resonance robot model with a series mechanical filter to restore passivity, as proposed in [35].

transducer is left out of this model.

This configuration was derived to permit large reductions in apparent inertia while preserving a passive interaction port. The mechanical filter augments a nonpassive port function and makes it passive at the new port on the environment side of the filter. This can be seen by modeling the mechanical filter as a two-port, like that shown in figure 3-3, that interacts with the force-feedback-controlled robot on one side and with the environment on the other. Figure 5-5 shows a schematic of a two-port between a robot with port impedance Z_{rob} (or admittance Y_{rob}) and the environment. If the two-port for this particular problem takes the form of a spring and damper in parallel, as shown in figure 5-6, its equations in impedance causality are:

$$Z_{filt} = \begin{bmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{bmatrix} = \begin{bmatrix} \frac{-b_s s - k_s}{s} & \frac{b_s s + k_s}{s} \\ \frac{-b_s s - k_s}{s} & \frac{b_s s + k_s}{s} \end{bmatrix} \tag{5.6}$$

Using this expression and the arrangement shown in figure 5-5 the admittance at the port that meets the environment, Y_2, can be derived:

$$Y_2 = \frac{1}{Z_{22}} + Y_{rob} = \frac{s}{b_s s + k_s} + Y_{rob} \tag{5.7}$$

The net effect of the series elements is to add a term, $\frac{s}{b_s s + k_s}$, to the port admittance. Recall that passive systems have port functions with positive real parts at all
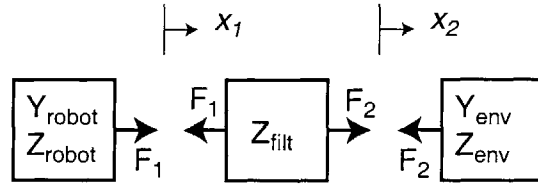
Figure 5-5: Generic two-port mechanical filter placed between the interaction port of a robot and the environment.
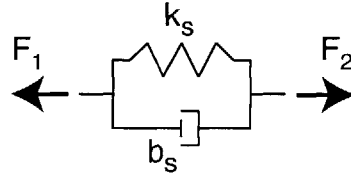
Figure 5-6: Parallel spring-damper as a specific two-port example for the generic model in figure 5-5.

frequencies. The real part of the term $\frac{s}{b_s s + k_s}$ when $s = j\omega$ is:

$$\Re\{Y_{filt}(j\omega)\} = \frac{b_s \omega^2}{b^2_s \omega^2 + k^2_s} \tag{5.8}$$

This term is greater than zero for all ω, so it makes the port more "positive real." It approaches $\frac{1}{b_s}$ at high frequencies, and is greater than $\frac{1}{2b_s}$ at all frequencies above $\omega = \frac{k_s}{b_s}$.

The strategy that Dohring et al. recommend for designing a mechanical filter is to first design a controller that delivers the desired impedance/admittance properties at the robot port Y_{rob}, regardless of passivity or coupled stability properties. The next step is to evaluate (by model or experiment) the real part of the controlled robot port admittance, to determine its minimum value $R < 0$ (if it is nonpassive) and the minimum frequency at which it is nonpassive, ω_r. Then b_s is selected such that:

$$\frac{1}{2b_s} > -R \tag{5.9}$$

209

and k_s such that:

$$\frac{k_s}{b_s} \leq \omega_r \qquad (5.10)$$

The role of the damper b_s is key, as it drains the excess energy produced by the nonpassive port. The spring k_s must be soft enough such that it permits b_s to drain this energy in the correct frequency range.

Because the mechanical filter is designed subsequent to controller design, it is important that the filter be made as stiff as possible. This is because the robot is designed to have the desired stiffness and other dynamic properties without the filter. The softer the filter, the more the port dynamics deviate from their target values. Despite this weakness, this technique provides a simple way to design series dynamics even if no model of the robot is available. The design can be based on data (the measured port admittance) rather than a model, and this technique can make any robot port passive given enough information. It is noteworthy that a parallel spring and damper attached to any controlled system port can make the port passive, with lower damping and stiffness required to passivate ports that are more highly non-passive to begin with.

Nonzero stiffness k_s is not necessary in the design of this type of mechanical filter. A series damper alone can passivate any system. This could be useful if unlimited travel is available at both the actuator and the environment, as in the case of rotary actuation. A highly geared electric motor with aggressive force feedback could be passivated by adding a series damper, possibly creating a high-quality force source. On the other hand, if the series damping is too low it drains a great deal of energy (this is how it passivates the system), so such an actuator would be highly power-inefficient. This may provide an opportune tradeoff between power-efficiency and force density. This is an interesting area that warrants further study in subsequent work.

5.2 Integrated design of series dynamics and control

While both design approaches presented above are reasonable, neither provides a systematic means for designing the series dynamics and controller together to facilitate their simultaneous optimization. Both require independent design of series elements and control. As a result, in both cases the series dynamics compromise a performance objective, and their own stiffness must be maximized subject to the constraints that allow them to stabilize contact. An oversimplified interpretation of these two design strategies is that the series elastic approach selects k_s only to change the physical structure of the system and change performance (while reaping the coupled stability benefits it provides, but with only empirical assurances of stability), while the mechanical filter approach selects k_s and b_s exclusively on the basis of stability (or passivity) with no recourse if performance (port behavior) is compromised too greatly.

A stiff series transmission is undoubtedly a worthy goal for many applications (for example those that require high bandwidth or strong position coupling), but in other cases softer dynamics may not be a detriment, provided that the controller compensates for them. The approaches described above prove that placing dynamics in series can help to reduce port impedance and stabilize contact, but perhaps because they do not focus on human interaction, they do not capitalize on the complete capability of series dynamics to "shape" the port impedance for the best performance. Because target performance is spring-damper behavior, placing a spring-damper *impedance shaper* on the robot is likely to help significantly, especially if the dynamics of the impedance shaper approach target dynamics. The term "impedance shaper" expresses the fact that the location of series dynamics is ideal for sculpting them to play an important role in endpoint impedance, not just to transmit forces. Properly shaping the series dynamics can tremendously reduce the error in port behavior. This point distinguishes this application from the prior work, and is explored in this section.

An effective systematic method would manage the tradeoff between stiff and soft

elasticity and evaluate performance against some ideal. It would allow concurrent design of control and any new dynamics. And it would permit alternative measures of stability, such as complementary stability. A new design method is presented below.

5.2.1 Applying search method

The search method used in the previous chapter to design force feedback controllers can easily be adapted to include mechanical design parameters. Provided that the structure is given, the method can be applied to search a combination of mechanical and control parameters to design an impedance shaper along with appropriate control to optimize performance. Rather than fighting the tradeoff that softer dynamics improve coupled stability but compromise performance, this method is capable of rapidly searching many combinations to find the optimal one(s). This approach can also include information about the environment, if available, to reduce conservatism.

5.2.2 Purely mechanical benefits

As explained and shown by example in section 5.1.1, for this application there are benefits to adding an impedance shaper independent of any force-based control. If the target system stiffness and damping are K and B respectively, a simple impedance controller applied to the model in figure 5-4 takes the form:

$$F_a = -K_2 x_1 - B \dot{x}_1 \qquad (5.11)$$

To implement the correct stiffness, K_2 can be set according to equation 5.4. Because the motion feedback is collocated with the actuator, this system is passive for all non-negative gains and therefore stable when coupled to any passive environment.

For a fixed target impedance defined by K and B, the search algorithm can be used to vary k_s and b_s and evaluate the cost, defined as in chapter 4 in terms of the port impedance magnitude. K_2 is varied with k_s according to the constraint equation 5.4. Here the relevant impedance is $Z_p = \frac{F_e}{x_3}$. Only stiffness values for k_s larger than K are considered, as $k_s < K$ would make it impossible for the system to satisfy the
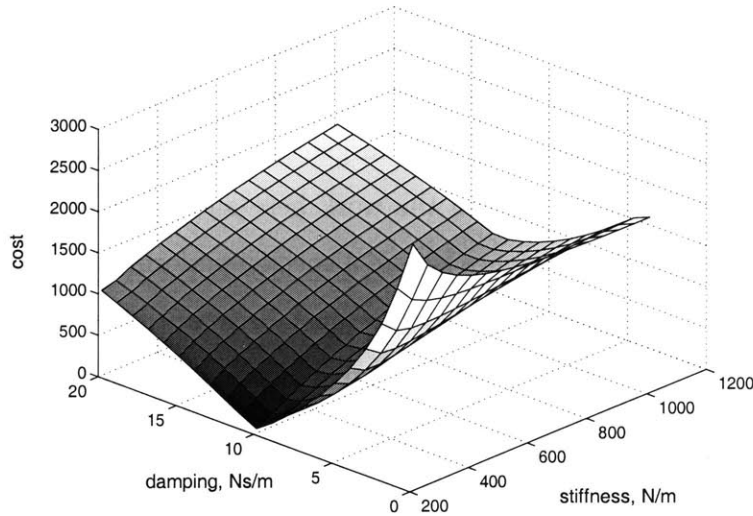
Figure 5-7: Cost versus series stiffness and damping. Simple impedance control on robot. Target stiffness 200 N/m, target damping 10 Ns/m

low-frequency target, and $k_s = K$ would require infinite K_2.

Figure 5-7 shows the cost versus k_s and b_s with a desired stiffness $K = 200$ N/m and desired damping $B = 10$ Ns/m. The cost is minimized when k_s approaches K and b_s approaches B. This result is neither surprising nor does it require an algorithm of such complexity to divine. The result is important, however, and worthy of brief attention. Because there is no mechanism to reject unwanted damping or inertia (i.e. force feedback), the design that minimizes these terms most effectively derives all or most of the target impedance characteristics from the impedance shaper (k_s and b_s) by stiffening the motion feedback gains on the inertial plant such that it resembles a pure motion source. In the extreme case, the inertia is servoed along the virtual trajectory, while the impedance shaper provides the nodic impedance, purely in hardware. This design is attractive for several reasons. No force feedback whatsoever is used, so nominally the system should be perfectly stable and perfectly passive. Also, the high-impedance properties of the inertia become irrelevant.

Practically, this design is extremely limited because it is only capable of accurately implementing a single impedance. Designing systems for multiple operating

213

conditions is discussed below. However, the result is interesting in its own right. In contrast to earlier work, which finds that the optimal series stiffness is the highest that provides stability or the lowest that satisfies performance requirements, here the optimal stiffness is found to be the target endpoint stiffness. This conclusion results from two differences between the work here and the literature. The first is that matching the target impedance is here a specific performance objective, and the second is that the controller stiffness can adjust to compensate for the presence of the series stiffness. Making the physical system closer to the target system eases the demands on the control, improving performance without sacrificing stability properties at all. This idea is extended to more practical control architectures in the rest of this section.

5.2.3 Mechanical filters for passivity with fixed nonpassive control

To connect this work to the prior art and to verify both approaches, the search algorithm was used to search for spring and damper parameter combinations that would passivate a robot port that is nonpassive due to a fixed feedback controller. This is the same problem addressed by Dohring's work and described in section 5.1.3. An impedance controller with proportional force feedback gain $K_{dc} = 10$ is imposed on the model in figure 5-4 with the control law:

$$F_e = -K_2 x_1 - B\dot{x} + K_{dc}(F_e - K_2 x_1 - B\dot{x}) \tag{5.12}$$

In this case, the control is not altered as the series dynamics are changed, so $K_2 = K = 200$, the target stiffness. The target damping is $B = 0$, the model is the robot in figure 5-4, and the rest of the parameter values are given in table 4.1. With this control applied, the real part of the port admittance is negative for frequencies above 581 rad/sec, with a minimum value of -0.32. From equations 5.9 and 5.10, Dohring's approach dictates damping $b_s \le 1.56$, and stiffness $k_s \le 581 b_s$.

Applying the design algorithm and searching over stiffness parameters from 110
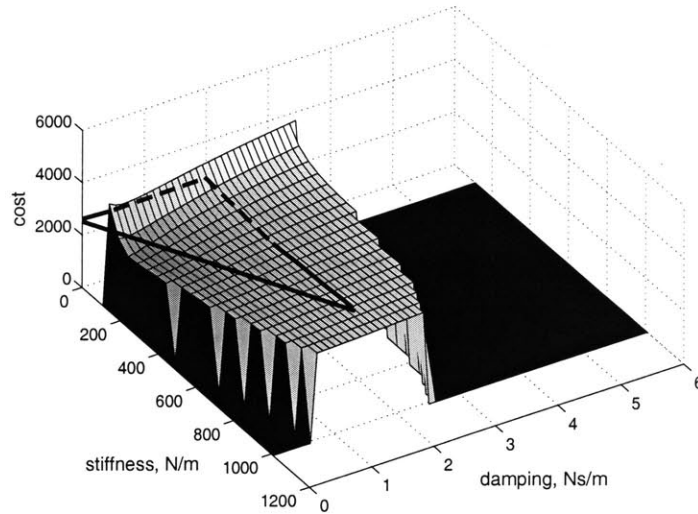
Figure 5-8: Cost versus series filter stiffness and damping for a fixed nonpassive controller. If stiffness-damping combination does not passivate, cost is set to zero.

to 1010 N/m in increments of 50 and damping parameters from 0 to 6 Ns/m in increments of 0.1 produces the plot of cost versus k_s and b_s in figure 5-8. Passivity is evaluated as the stability measure, in lieu of nominal and robust stability. Combinations of specific parameters for k_s and b_s that do not passivate the interaction port are unacceptable, and are shown as zero cost. All of the solutions identified by Dohring's criteria are duplicated here, and are indicated by the interior of the heavy triangle in the figure. However, the set of viable solutions is also expanded considerably. This is because Dohring's method is conservative; the parameters are selected such that the filter adds positive real admittance equal to or greater in magnitude than the maximum negative real admittance at the port, across all nonpassive frequencies. The search method admits solutions that take advantage of the shape of the mechanical filter's real admittance curve to add sufficient positive real admittance in the proper frequencies. Of course, some conservatism may be wise in order to accommodate modeling uncertainties and other errors. Regardless, the search method finds results that are comparable to Dohring's predictions, with differences explainable by the differences in the methods.

215
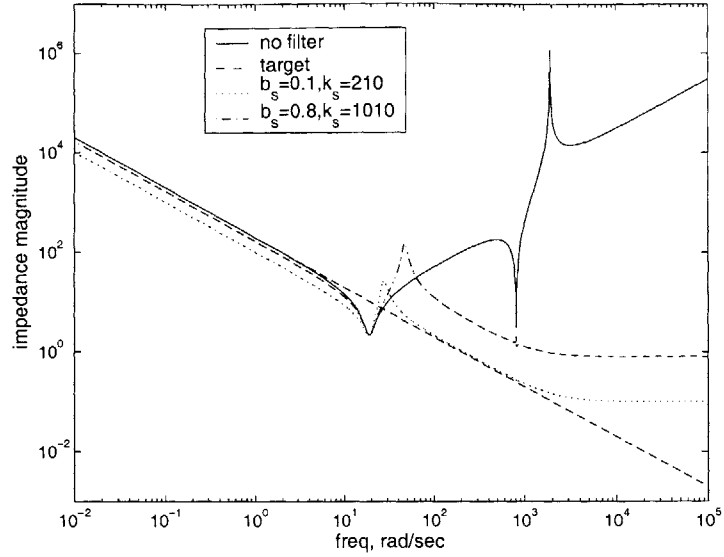
Figure 5-9: Robot port impedance magnitude with no mechanical filter and with two different mechanical filters that passivate the interaction port. Control is fixed, with $K_{dc} = 10$. Target impedance is also shown.

Additionally, the search method provides an alternative means to select from among the passivating mechanical filters (or impedance shapers). Among the passivating solutions, the cost is lowest (2414) when the stiffness is 210 N/m (the closest tested value to the target of 200 N/m) and the damping is 0.1 Ns/m. This is somewhat surprising, as stiffer filters are expected to better preserve the target impedance as implemented by the controller. To understand this, the impedance magnitude of the controlled system with no series dynamics and with two different sets of passivating springs and dampers are plotted along with the target impedance in figure 5-9. The lowest-cost solution, $k_s = 210$ and $b_s = 0.1$, is shown with a stiffer solution, $k_s = 1010$ and $b_s = 0.8$, that also passivates the system. While the stiffer dynamics lead to better performance at the low frequencies where stiffness dominates, at high frequencies the softer filter performs better. Between 100 and 1000 rad/sec, the stiffness k_s dominates; with the softer filter, this is closer to target behavior. At frequencies higher than 1000 rad/sec, the filter damping b_s dominates, and again the softer filter performs better. If matching stiffness at low frequencies is a higher priority than behavior beyond resonance, the cost can be altered or weighted accordingly.

216

Improving the performance measure

As in the nonzero stiffness analysis from the previous chapter, the unweighted cost does not make a perfect performance measure. Here this is even more problematic because the inertia is not important beyond its resonant frequency with the mechanical filter dynamics, and the reduction in inertia only appears as a shift in this resonant frequency (upward as inertia is reduced). This characteristic is telling, because it shows the reduced influence of the inertia on the endpoint impedance. Because the motion of the mass is decoupled from port motion, inertia is far less important than it is when no impedance shaper is used. Unless the series dynamics implement the nodic impedance exactly and the performance goal for the robot is motion control, the inertia still contributes to endpoint impedance. There are several ways the cost can be modified to capture the inertia more accurately. One approach is to include a weighting function in the cost that decreases with frequency so that low-frequency impedance errors are penalized more than high-frequency errors. Since errors occur at the resonant frequency, this change punishes low-frequency resonance (and high inertia) and reward high-frequency resonance (and low inertia). Alternatively, the effective force gain magnitude in particular frequency ranges could be directly included in the cost, to encourage higher gains. It is also important to include the effective rejection of low-frequency nonlinear friction that is present in the actual system but not in the idealized model. Including force feedback gains at certain frequencies in the cost is one approach to this, though this strategy must include consideration of the relative motion of x_2 and x_3. If k_s is much softer than K_2, then little motion of x_2 is required as x_3 moves, and friction at x_2 has little effect. To the contrary, if k_s is much stiffer than K_2, greater motion of x_2 is required as the port x_3 moves, and friction becomes very important. This difference must be included if the cost is to accurately reflect the degree to which unmodeled friction is compensated.

It is also possible to include stability (or passivity) margins in the cost. The ease of determining an appropriate quantity to measure the margin depends on the particular problem. In general, this is not as straightforward as it is in classical servo control

design, because of the structural differences cited in chapter 3. Changes to control do not always have obvious implications in the transfer function that determines stability. However, in many examples there is at least one parameter, often a controller gain, with a straightforward relationship to system stability. A stability margin can be defined with such a parameter.

For the examples described in this thesis, the standard unweighted cost of equation 3.19 suffices to demonstrate the differences between various combinations of hardware and control. Different problems may require the measures suggested above.

5.2.4 Complementary stability and fixed nonpassive control

An important expansion of the idea of impedance shaping hardware is to the more appropriate, environment-based complementary stability measure described in the two preceding chapters. If certain dynamics in series with a nonpassive robot can render the output passive, it stands to reason that similar dynamics may permit even greater improvements in performance if passivity is relaxed. When passivity is replaced by complementary stability, the set of acceptable impedance shaper parameters is expected to expand, providing greater flexibility. If the control is allowed to vary, higher gains are expected. The examples in the previous chapter demonstrate that performance can be enhanced when passivity is relaxed, and that conclusion may apply here as well.

The preceding example was repeated using complementary stability (with the same environment used throughout chapter 4) instead of passivity as a solution constraint. The robot, control, and environment parameters are the same as used in the example in chapter 4. In that example, it was shown that the robot (with no series dynamics) only satisfies complementary stability with the environments for proportional force feedback gains $K_{dc} \leq 1.74$. Since the gain used here is $K_{dc} = 10$, this condition is not satisfied. Thus without any series dynamics, the robot is unstable when coupled to at least one environment in the design set. Figure 5-10 shows the set of solutions produced by this algorithm run. A comparison of figures 5-8 and 5-10 reveals three differences between using complementary stability and passivity

as the stability constraint. First, the complementary case admits $b_s = 0$. This has interesting implications. While Dohring shows that damping is essential to passivate the port, Pratt's group has had success in real applications without damping explicitly designed. This analysis suggests a reason for this discrepancy, and it relates to the properties of the environment. Although the interaction port is nonpassive, the environments in the specified set do not induce instability of the coupled system. The unstable coupling between a robot and this class of environment is stabilized with the addition of an undamped spring in series.

A second difference between the passivity and complementary stability results is that complementary stability admits high-stiffness, low-damping filter solutions. This may be useful if the robot's controlled port impedance is to be transmitted as accurately as possible, without added damping losses. Finally, complementary stability admits higher stiffnesses at moderate damping values (between 2.5 and 3.1 Ns/m) than the passivity constraint.

As expected, complementary stability admits more impedance shaper combinations than passivity. The zero-damping case is especially important in practice, and some of the other cases may similarly be useful. The use of complementary stability in lieu of passivity is again predicted to have a significant impact on performance.

5.2.5   Variable control with series dynamics

A far more dramatic improvement can be achieved when the control is allowed to vary along with the impedance shaper parameters. In the previous two examples, soft dynamics cause deviation from the target stiffness because the controllers implement this stiffness exactly on the robot, so that adding additional dynamics reduces the stiffness below the target. Furthermore, because the control is predetermined, there is no reward (in the system cost) for the impedance shaper that passivates or stabilizes the system more than others; the system is not able to take advantage of excess stabilizing capability to reduce unwanted impedance more aggressively.

To remedy these issues, two control parameters are allowed to vary. The first is the controller spring stiffness K_2, which is constrained to the target stiffness K

219
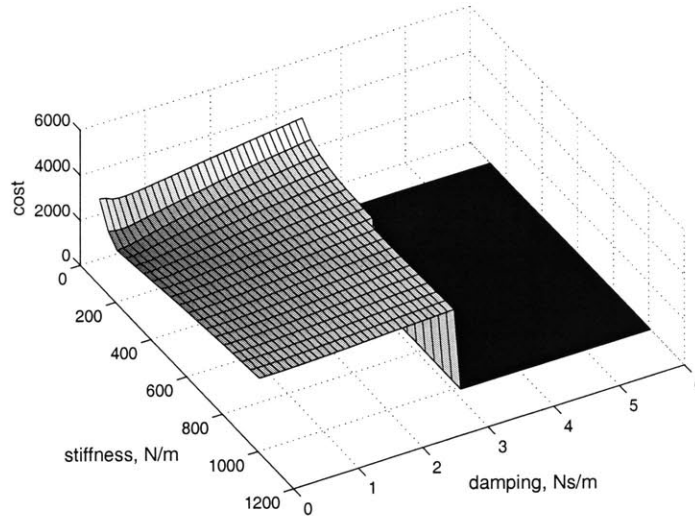
Figure 5-10: Cost versus series stiffness and damping. Robot uses proportional force feedback control with fixed gain, $K_{dc} = 10$. Solutions that satisfy complementary stability are plotted; solutions that do not satisfy complementary stability are shown with $C=0$.

and the impedance shaper stiffness k_s by equation 5.4, to ensure that the stiffness is implemented accurately regardless of series dynamics. The second variable control parameter is the proportional force feedback gain K_{dc}, which is varied independently. As compared to the results in the previous section, this change has several important effects. For impedance shaper parameter combinations that were ruled out in the previous example, this method permits lowering the force feedback gain such that the system satisfies the stability requirements. For parameter combinations that satisfy the stability requirements, the maximum permissible gain value is found, permitting more aggressive rejection of unwanted impedance without compromising coupled stability.

With a target of $k_{des} = 200$ and $b_{des} = 0$, the series stiffness k_s was varied from 210 to 1010 N/m (in increments of 50 N/m) and the series damping b_s was varied from 0 to 6 Ns/m (in increments of 0.1 Ns/m). For each combination, with K_2 constrained as described above, K_{dc} was raised to its maximum stable value. The resulting cost at the stability boundary is plotted versus k_s and b_s in figure 5-11. The
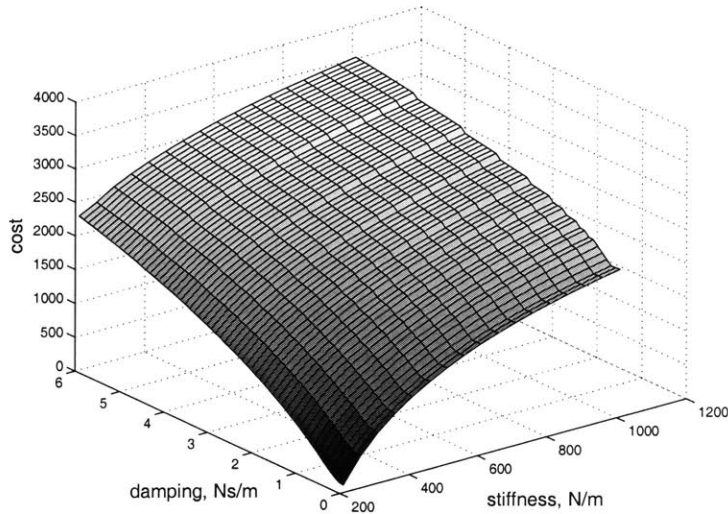
Figure 5-11: Cost at the stability boundary versus series stiffness and damping. Proportional force feedback gain raised or lowered to stability boundary for each combination of physical parameters.

improvement versus the fixed-control results is dramatic. While the minimum cost for the impedance shapers tested with fixed control is over 2000, here the cost is as low as 109. The advantage comes from the fact that the feedback gain can be increased above $K_{dc} = 10$ when the conditions permit. For $b_s \leq 0.2$, the gain can be raised to the maximum tested value, $K_{dc} = 100$. The superior performance of the best solution, with $k_s = 210$, $b_s = 0$, and $K_{dc} = 100$ is shown along with an inferior solution, the target behavior, and the system with no impedance shaper and the same fixed gain used above in figure 5-12. The soft shaper matches the stiffness at low frequencies and approaches it at high. The increased force feedback gain suppresses the robot inertia and moves the first resonance from around 20 rad/sec to higher than 200 rad/sec. The magnitude of the robot structural resonance around 800 rad/sec is also dramatically suppressed, because the impedance shaper so closely approximates the target dynamics that little motion is required at the robot port. It is noteworthy that the best-performing impedance shaper has zero damping, a case that would be excluded if a passive port impedance were required. Complementary stability permits this solution, and performance is improved.
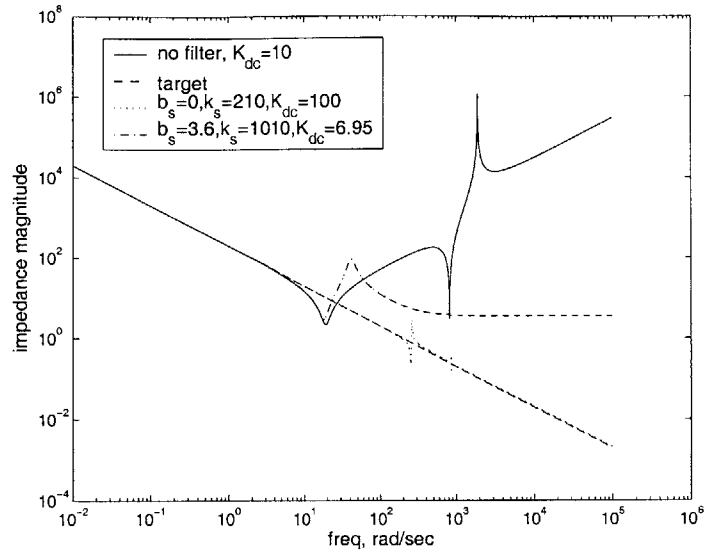
Figure 5-12: Robot port impedance magnitude with no series dynamics, with two parametrizations of series spring and damper, and target impedance.

Another interesting trend is evident from looking at the stable gain limits versus stiffness and damping, shown in figure 5-13. The gain increases steadily as the impedance shaper damping, b_s, is *decreased*, with the highest gains coming at zero or near-zero damping. Furthermore, there is no dependence whatsoever of the gain on stiffness (at least not within the discretization of the gain tested and the conservatism of this analysis). The improvements in cost as the *stiffness* is decreased come strictly from the fact that the physical system more closely resembles the target system, not from improved force feedback.

If reliable model information and access to the control algorithm are available, this approach is clearly superior to the previously described approaches. When control and design parameters are varied together, the full power of this method becomes clear. The previous approaches are included to introduce the concept, to relate it to the prior art, and to highlight the advantages of this approach applying the search method to both hardware and control parameters simultaneously. The methods with fixed control remain useful in the case that the robot port impedance can be determined only by experiment, and the controller is not accessible for modification.

222
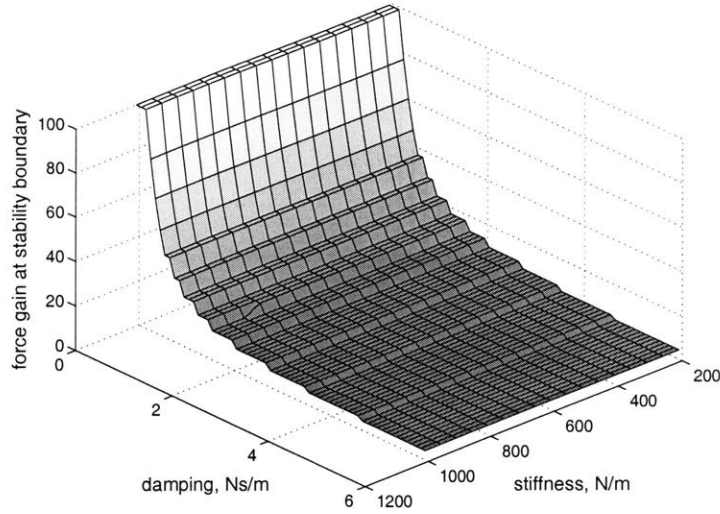
Figure 5-13: Proportional force feedback gain at stability boundary versus series stiffness and damping. In this example, the gain depends entirely on the damping.

Passivity

If passivity is required, control can similarly be varied to find the optimal value and configuration. The example from the previous section was repeated exactly, except that port passivity was used as the stability requirement instead of complementary stability. As a practical matter, making this change to the design algorithm is trivial and can speed computation by avoiding the calculation of structured singular values. Figure 5-14 shows the differences between the two cases. For each spring-damper combination, figure 5-14 plots the difference between the complementary stable gain limit $K_{dc,c}$ and the passive gain limit $K_{dc,p}$, divided by the passive gain limit. That is:

$$\frac{K_{dc,c} - K_{dc,p}}{K_{dc,p}} \tag{5.13}$$

For many of the stiffness and damping values, there is little or no difference between the two conditions. For the lowest damping values, however, and especially at high stiffnesses, the complementary stability measure permits much higher gains. In the extreme case, when the series damping is zero, passivity limits the force feedback gain to unity, but complementary stability permits gains of up to 100. From the analysis

223
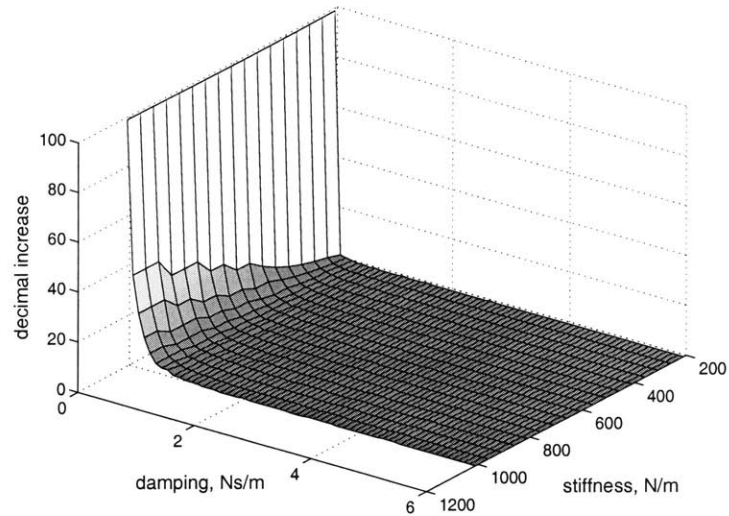
Figure 5-14: Difference between the maximum stable gain for complementary stability and for passivity, as defined by equation 5.13, versus series stiffness and damping. Differences are minimal except at low damping.

in section 5.1.3 it is clear why the high-stiffness, low-damping case does not passivate the impedance port; the positive real contribution of the mechanical filter is only at very high frequencies, beyond where it is useful. Thus the force gain must be kept so low that the robot is passive without a filter. It is more interesting, however, that this condition permits very high gains when complementary stability is the performance measure. It is possible that the effective environment stiffness is reduced enough by this series spring to render it less problematic. It is also possible that the benefit comes from decoupling the environment's mass from the robot port. Regardless, this configuration is promising.

5.2.6 Multiple operating conditions

The primary result found above, that the impedance shaper dynamics should approach the target dynamics, is somewhat unsatisfying and partially flawed. The extreme case of zero target impedance exposes this flaw. If the algorithm is used to search for the best series dynamics, and performance is based on minimizing

impedance alone, the solution tends toward zero stiffness and zero damping. This makes perfect sense given the way that the problem is defined, but this clearly does not produce a practical design, as force cannot be transmitted from actuator to environment if there is nothing but empty space in between.

There are several likely sources of practical bounds on series dynamic elements, particularly the stiffness. If force is to be transmitted at high frequencies to a kinematically fixed environment, low series stiffness means that the source actuator must move significantly to vary output force. Because the actuator mass must be accelerated, this limits bandwidth. If the source actuator and structure at the environment have limited travel (as for most linear actuation methods), then some minimum stiffness is required to maintain a reasonable positional coupling between the actuator and environment. This is needed to allow impedance to be rendered as desired over the entire travel without encountering limits of travel at either side. Similarly, almost all applications require some range of nonzero impedance to be rendered. Stiff impedance requirements cannot be satisfied without stiff dynamics physically in series, and the same physical system must generally be used for both stiff and soft (or zero) impedance. So, target impedance cannot be matched in hardware for the soft cases.

There are several ways to handle the need to select hardware that improves performance for multiple target impedance specifications. If the number of independent target impedance regimes is low, it may be practical to combine the optimizations for the multiple specifications. If a search is done across identical parameters for each target, the optimizations can be combined simply by adding the costs (the sum can be weighted if a certain specification is more critical). Once the individual optimizations have been run this step is trivial, as the algorithm returns an array containing the cost at the stability boundary for each parameter combination; these arrays must simply be summed to determine the composite cost. If n different impedances are targeted, and the search algorithm is used to compute a cost C_i for each as a function of m different (hardware or control) parameters P_j, a composite cost C_c can be determined

by the equation:

$$C_c(P_1, P_2, ..., P_m) = \sum_{i}^{n} W_i C_i(P_1, P_2, ..., P_m) \qquad (5.14)$$

W_i is a real scalar weight for the i^{th} cost and can be used if some operating conditions are more important than others. The minimum composite cost across the parameter combinations gives the least total error. The composite cost could be augmented to include other measures if needed; an example is the variance of the individual costs, which would be useful if spreading the error evenly over the several operating conditions is an objective. If a subset including k of the parameters P_j are allowed to vary across different impedance conditions (e.g. controller gains) while the remaining $m - k$ parameters (e.g. physical hardware parameters) must be common to all, equation 5.14 can be adjusted such that each cost C_i is replaced with a minimum cost:

$$C_{mc}(P_1, P_2, ..., P_{m-k}) = \sum_{i}^{n} W_i C_{mi}(P_1, P_2, ..., P_{m-k}) \qquad (5.15)$$

C_{mi}, the minimum cost function for the i^{th} impedance condition, is determined at each combination of its arguments by fixing $(P_1, P_2, ..., P_{m-k})$ and searching all combinations of the parameters $(P_{m-k+1}, P_{m-k+2}, ..., P_k)$ for the lowest cost. This can easily be done by post-processing after running the search method. The parameters $(P_1, P_2, ..., P_{m-k})$ are chosen by selecting the minimum value of C_{mc}. The values of $(P_1, P_2, ..., P_{m-k})$ are then selected individually for the i^{th} impedance condition by determining which values produce the entry of C_{mi} that corresponds to the parameters $(P_1, P_2, ..., P_{m-k})$. Two examples are provided below.

Example 5.1: Suppose that an actuator is to operate with three different port impedances: *a)* $k = 1000$ N/m, $b = 0$ Ns/m, *b)* $k = 500$ N/m, $b = 30$ Ns/m, and *c)* $k = 0$ N/m, $b = 20$ Ns/m. The robot and environment models and the cost function are the same as used in this and the previous chapters. A spring-damper impedance shaper characterized by stiffness $k_s = P_1$ and damping $b_s = P_2$ is to be placed in series between the

226

source actuator and the environment, and a proportional force feedback impedance control law is to be used that is characterized by the gain $K_{dc} = P_3$. A single set of optimal hardware parameters P_1 and P_2 is to be found for all three cases, but the control parameter P_3 is permitted to vary across endpoint impedance specifications. Thus the functions C_{mi} have P_1 and P_2 as their arguments, and can be represented by surface plots. Figure 5-15, panels a), b), and c), shows the minimum cost function versus k_s and b_s for each of the three endpoint impedance specifications. In each case, the function shows the minimum cost across all stable values of K_{dc} for each combination of k_s and b_s. Stiffness values from 1100 to 2500 N/m were considered (such that the device stiffness always exceeds the target stiffness), along with damping values from 0 to 50 Ns/m. Figure 5-15, panel d) shows C_{mc}, the sum of the three cost functions for the various operating conditions. From this plot, the lowest composite cost is achieved when $k_s = 1100$ N/m and $b_s = 20$ Ns/m. This is also the lowest cost for condition c), while the errors that it introduces for conditions a) and b) are relatively small.

Example 5.2: In this example, the robot system of the previous example is to exhibit two different target impedances, the first with target stiffness $k = 1000$ N/m and target damping $b = 0$ Ns/m, and the second with target stiffness $k = 1000$ N/m and target damping $b = 40$ Ns/m. The individual minimum cost functions are plotted versus k_s and b_s in panels a) and b), respectively, of figure 5-16. In this case, the composite cost includes the sum of the individual functions as well as a fraction of the difference between the two. This tends to select solutions where the error is similar between the two cases. Specifically, $C_{mc}(k_s, b_s)$, as defined by the equation below, is plotted in figure 5-16, panel c):

$$C_{mc}(k_s, b_s) = C_{m1}(k_s, b_s) + C_{m2}(k_s, b_s) + 0.3|C_{m1}(k_s, b_s) - C_{m2}(k_s, b_s)|$$

$$(5.16)$$

227
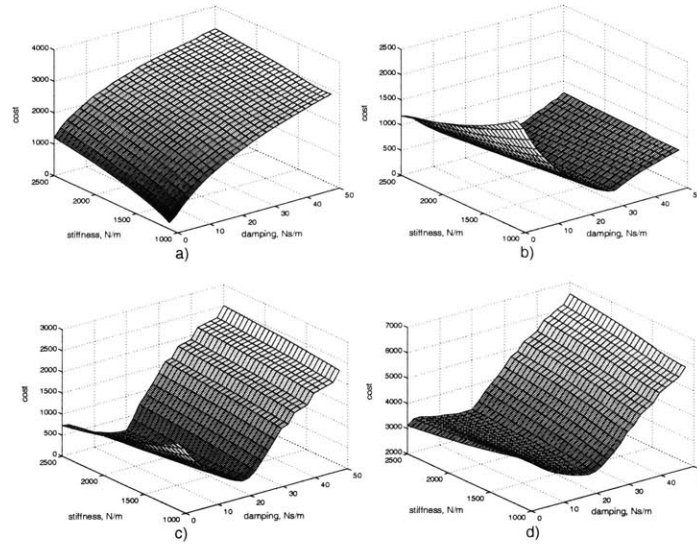
Figure 5-15: Cost versus k_s and b_s with a target impedance of *a)* $k = 1000$ N/m and $b = 0$ Ns/m, *b)* $k = 500$ N/m and $b = 30$ Ns/m, and *c)* $k = 0$ N/m and $b = 20$ Ns/m, and *d)* composite cost given by the sum of the three surfaces in *a)-c)*.

In this case the best solution features $k_s = 1100$ N/m and $b_s = 14$ Ns/m, a damping value between the two target impedances. Another solution that has cost nearly as low is with $k_s = 2500$ N/m and $b_s = 0$ Ns/m, featuring high stiffness so that the apparent damping is modulated primarily by the controller, not the impedance shaper.

Choosing impedance shaper parameters for multiple operating conditions may not require computation of a composite cost. The basic trends borne out in the results presented here can be used to develop some rules of thumb to guide design. For instance, no spring with stiffness less than the maximum target stiffness should be considered, because this would preclude the possibility of achieving the target exactly. If all the other operating conditions use stiffness that is significantly below the maximum, then the impedance shaper damping can probably be designed to approach the high-stiffness operating point, because the robot control will dominate motion at lower target stiffnesses. On the other hand, if more than one operating condition with stiffness close to the maximum have radically different damping, it is likely that the impedance shaper stiffness will have to be considerably higher than
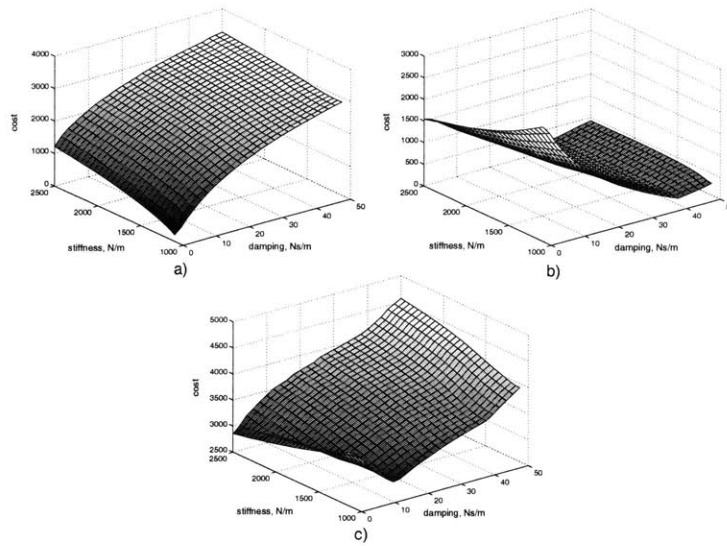
Figure 5-16: Cost versus k_s and b_s with a target impedance of *a)*$k = 1000$ N/m and $b = 0$ Ns/m and *b)*$k = 1000$ N/m and $b = 40$ Ns/m, and *c)* composite cost given by equation 5.16.

the maximum target stiffness, so that the controller dominates damping.

5.2.7 Alternative parameter variation

The parameter search design method presented here can be used in a variety of ways depending on which parameters are fixed and which are variable. This is true both for control parameters and hardware parameters. Several cases have been presented in the examples in this and the preceding chapter, including: fixing the hardware design while varying the control, fixing the control while varying the impedance shaper parameters, and varying several hardware and control parameters together. The method is flexible in that any combination of parameters can be varied, as long as the structure is provided. It is convenient if one parameter, such as a gain, is monotonically related to the cost, but this is not necessary. The examples detailed here illustrate common cases, but represent only a small subset of possible applications of this approach.
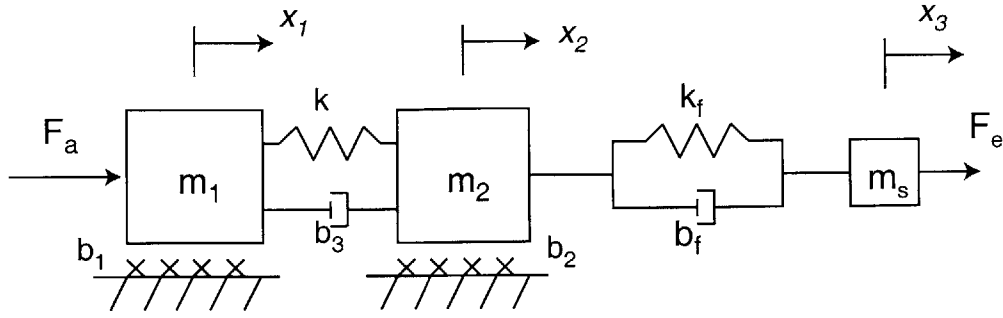
Figure 5-17: Mass m_s placed in series between the robot and the environment. The spring k_f and damper b_f are due to the force transducer.

5.2.8 Inertial series dynamics

In addition to the proposed arrangement of "soft" dynamics (springs and dampers) in series with the environment, alternative dynamic elements can also have advantages. Spring-like environments can provoke instability when coupled to systems that use force feedback. When the environment has inertial properties, these problems are often not as acute. This property can be exploited by adding inertia to the robot, so that all environments it encounters appear to have this inertial characteristic.

Because a primary objective is to reduce apparent inertia, *adding* physical inertia is a counter-intuitive solution. A critical factor is the placement of the inertia. It must be on the environment side of the force sensor, so that to the robot controller it appears to be a part of the environment rather than the robot. A mass in this configuration with the single-resonance robot model used throughout this thesis is shown in figure 5-17. This model includes the stiff (fixed) force transducer dynamics, represented by k_f and b_f. Force is measured across these elements, between x_2 and x_3, the motion of m_s.

The environment was assumed to have stiffness k_h and damping b_h with the range of values specified in table 4.1. The environment mass m_h was assumed to be negligibly small (several orders of magnitude smaller than m_s), with a nominal value $m_{hn} = 10^{-8}$ and a maximum deviation of $m_{hd} = 10^{-9}$. A proportional force feedback law with a target impedance of zero was used, as in equation 4.15. The search al-

gorithm was run with m_s and K_{dc} varying. For each value of m_s, K_{dc} was raised to its maximum value that met the coupled stability condition. As K_{dc} is increased, the apparent friction and mass due to the elements between the actuator and sensor are reduced. In particular, the apparent inertia due to m_1 and m_2 is reduced by the factor $(1 + K_{dc})$. The total endpoint inertia (at frequencies below the robot's structural resonance) is a sum of the series mass m_s and the apparent mass due to m_1 and m_2:

$$m_{app} = \frac{m_1 + m_2}{1 + K_{dc}} + m_s \tag{5.17}$$

Figure 5-18 shows the gain at the stability boundary K_{dc} and the resulting apparent mass m_{app} versus m_s. The resulting apparent Coulomb friction, reduced from a physical value of 20 N, is also shown. As m_s is increased, steadily higher force gains are permitted without destabilizing the coupled system. The increase in K_{dc} is generally sufficient to offset the increase in apparent inertia due to the added mass m_s, such that m_{app} stays nearly constant regardless of m_s, and the apparent friction decreases substantially. As 1 kg of mass m_s is added, the apparent mass m_{app} increases by less than 0.2 kg, less than 20% of the added physical mass. However, the gain K_{dc} doubles from 2.3 to 4.6, rejecting friction 70% more effectively, and the total output impedance is reduced. Thus adding inertia to the system can actually help to reduce the apparent impedance.

Experimental verification

This finding was tested on the screw-driven vertical module described throughout this thesis. Two different types of controllers were used: a proportional force feedback controller and a lag compensator like that designed in chapter 4. A single spring environment with stiffness of $16,400$ N/m was used for all stability testing. Both controllers used a virtual spring with 50 N/m stiffness to induce contact with the environment. The handle was removed from the robot end-effector to minimize the effective environment inertia. First the coupled stability gain limit was identified for both controllers. This was done by varying the DC force gain and testing stability

231
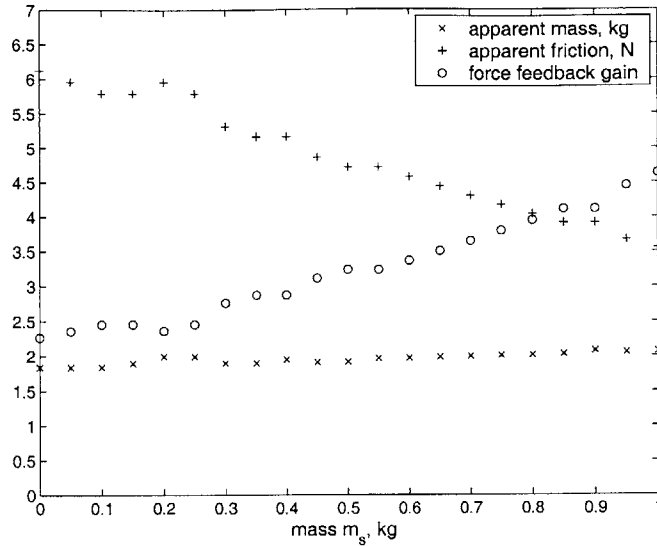
Figure 5-18: Apparent mass, friction, and proportional force feedback gain at complementary stability boundary versus series mass, as predicted by design algorithm.

as described in section 4.4.1. The "impact" test was used; the controller was judged stable only if it could preserve stability through an impact with the spring, and decrease oscillations to rest. Then a 500 g mass was bolted to the force transducer and the test was repeated, identifying a new set of stable gain limits.

With proportional control, the bare system (without the mass) coupled stably to the spring at gains up to 1.9. When the mass was added, stable coupling was achieved at gains up to 2.5. Assuming a 6 kg physical robot inertia, this means that the apparent mass at the stability boundary is initially 2.07 kg (with force feedback gain of 1.9). When the 500 g mass is added, the apparent mass at the stability boundary increases by 140 g to 2.21 kg. Similarly, with lag control, the bare system coupled stably to the springs with DC gains up to 2300. With the added endpoint mass, the gain could be increased to 3300. From the model of the system, this translates to an increase in apparent mass from 1.25 kg to 1.37 kg In both cases, the added mass permits an increase in force gain that offsets more than 70% of the extra mass. This is close to the analysis prediction that increased force gain could offset as much as 80% of the added mass, in a limited range.

232

As demonstrated in the previous chapter for control design, the search method requires the assumption of structure and searches for the best parameters. When this process is repeated for several different structures, general trends emerge that suggest which approaches are most valuable. An analogous procedure can be used to design physical impedance shapers in conjunction with control laws that improve the rendering of target impedance. The results here suggest that inertia, compliance, and damping can each have value. Several runs of the search method may reveal which approach is most effective for an application. Other factors related to implementation may also play an important role in design. For instance, a particular practical implementation may make it easy to add inertia, but difficult to add compliance. This algorithm can determine whether adding inertia improves or degrades the port impedance. The following chapters discuss one possible means of implementation.

5.3   Series dynamic conclusions

Adding series dynamics can improve interaction both by altering the actuator structure to make it more closely resemble the target system, and by permitting more aggressive force feedback to reduce unwanted actuator impedance. While several investigators have proposed series springs or series spring-dampers to improve interaction, this work contributes several new results to interpret some of the earlier work and to guide design.

Unlike in the prior art, here it has been shown that with a high-impedance source actuator, the ideal series dynamics approach the target dynamics (particularly when spring-damper behavior is desired). This serves to reduce the error due to the intrinsic differences between the hardware and the desired behavior, and also to make the system more amenable to force feedback.

By applying the control design technique developed in the previous chapters, a systematic design method for series dynamics has been presented. Given a physical system structure, this permits the parameters to be chosen on the basis of optimizing performance. Design can be performed either with fixed control, which resembles

work previously reported in the literature, or with the control varying with the other design parameters. This last approach, newly proposed here, permits dramatically better performance as measured by the endpoint impedance.

The design method also permits the simultaneous optimization of one physical system for multiple operating conditions. Once the system is in place, the hardware parameters stay constant but the control changes. This technique can ensure that performance is met for more than one operating condition.

Finally, it has been shown that adding inertial dynamics to a robot can help to ameliorate interaction stability problems in certain applications. This can permit more aggressive force feedback such that the added mass only slightly increases the apparent mass, while friction is reduced. This is a non-intuitive solution: although low endpoint impedance is desired, and inertia contributes to the endpoint impedance, *adding* inertia to the system can actually produce lower total endpoint impedance.

This chapter has explored potential changes to the physical structure of actuators to improve performance. In the next two chapters, remote transmissions are considered as a method to directly reduce the effective actuator weight. Their relationship to the series dynamics discussed in this chapter is explored, and the possibility of using the remote transmission to introduce desired dynamics is discussed in appendix B.

Chapter 6

Remote Transmission for Actuator Weight Reduction

The preceding chapters have described several methods for improving the ability of high-impedance actuators to render impedance. This can indirectly improve actuator force density, by permitting the use of actuators with high impedance that tend to have better force density than those with low intrinsic impedance.

This and the following chapter focus on using remote transmissions to facilitate more direct reduction in endpoint actuator weight. This approach could permit the source actuator to remain stationary even as the endpoint package moves. Specifically, a hydraulic remote transmission is proposed for use with various source actuators. In addition to transmitting power, this transmission can incorporate impedance shaping to also exploit the benefits described in chapter 5. In this chapter the basic configuration is introduced, and design calculations are described. In chapter 7 the design and testing of a specific prototype are detailed.

6.1 Hydraulic transmissions with electromagnetic actuators

Prior work on series elasticity and mechanical filters used mechanical elements to implement series dynamics with either electromagnetic [35, 95] or hydraulic [99] source actuators. The electromagnetic actuators were characterized by high impedance because they were highly geared. The hydraulic actuator used flow control and therefore also had high impedance.

Mechanical springs are convenient series elements because of their linearity and familiarity. However, other implementations are possible. Hydraulic systems have predictable damping and inertia, and can be designed to exhibit compliance and more complex dynamics. Their behavior is devoid of discontinuities like backlash and static friction, so they have an intrinsic appeal for human interaction. They also have a key feature shared with pneumatic systems that is not found in mechanical spring implementations: they can be contained in flexible, continuously deformable structures (i.e. hoses) without significant loss of performance. This presents the possibility of a new method for improving actuator force density. As shown in chapter 2, when an actuator is used in a serial configuration, its mass and weight increase both the load of the other robot actuators and the endpoint impedance. Chapters 3, 4, and 5 have focused on indirectly reducing actuator mass by permitting higher impedance and reducing its appearance at the endpoint by force feedback. If, on the other hand, the actuator can remain stationary as the robot end-effector moves, it does not contribute additional load. This is the idea behind using remote transmissions, also discussed in chapter 2. Fluid systems with flexible hoses may provide the opportunity to keep a source actuator stationary as the end-effector moves by transmitting the actuator force or torque to a package at the endpoint. While the endpoint package mass is critical, the source actuator mass is less important because it does not contribute to loading. Furthermore, the fluid transmission could change the actuator structure in advantageous ways identified in chapter 5.

This observation suggests at least two distinct but related architectures. The first

uses a stationary source actuator with high force capacity and low impedance but possibly large mass and a contained fluid volume to transmit actuator behavior to a low-mass end-effector. The second uses a moving or stationary source actuator with high impedance and force feedback, and a contained fluid volume to transmit forces and motion and shape the actuator port impedance to aid interaction, as designed in chapter 5. The fluid transmissions in these two cases have slightly different requirements. These are explained in the following section.

6.2 Requirements for two series fluid architectures

6.2.1 Low source impedance, transparent transmission

If an actuator can remain stationary (for example at the robot base), its weight is less important (minimizing total weight remains a goal for purposes of transportation). As shown in figure 2-10, electromagnetic actuators can have high force capacity with low impedance when their mass is large. On the other hand, servo-hydraulic actuators have excellent force density, thanks to the properties of fluids, but have high impedance that originates in their valving and control system. Here it is proposed to exploit the electromagnetic and hydraulic domains for their best features. Electromagnetics are used to generate force and represent impedance, but the heavy magnets, coils, and back-iron are kept stationary to avoid contributing to endpoint impedance. Fluids are used for their force transmission density, but the system is designed without small orifices through which the fluid must be forced.

The transmission can be represented as a two-port as in figure 5-5. The goal for the fluid transmission is *transparency*, to (somewhat liberally) borrow a term from the teleoperation literature [78].

> **Definition:** A transmission is defined as perfectly *transparent* if the same time history of force, motion, and energy is exchanged between the source actuator and the transmission (port 1 in figure 5-5) as between the transmission and the environment (port 2). The impedance at port 2 should

237

equal the actuator impedance, $Z_2 = Z_{rob}$.

To be transparent, the transmission should provide infinite stiffness with zero inertia and damping, with no transmission delay. Any actual system will have nonzero compliance, damping, inertia, and transmission delay. Some compliance can be tolerated, and may even be advantageous, as shown in chapter 5. Nonzero inertia and damping contribute directly to the endpoint impedance, as shown below; these are likely to be advantageous only if force feedback is used at the source actuator. Key to determining whether this design is feasible for a particular application, and to designing the specific system, is understanding the tradeoffs between geometric parameters and impedance properties. These issues are addressed at length below.

This configuration may permit extremely simple control such as simple impedance control. Feedback can thus be limited to signals from the source actuator alone. Because of this, the transmission dynamics are only important in their impact on the endpoint impedance. Although the objectives of high stiffness with low damping and inertia are common for transmission design, this is an important distinction. For example, the transmission should be stiff, but practically it must only be stiff enough to transmit the desired endpoint stiffness. An acceptable stiffness for this task may be much lower than if a control loop is closed around the transmission. A similar statement can be made about transmission delays, which have been shown to cause significant problems in physically similar systems [113].

This proposed system is simple, consisting of only an electromagnetic source actuator with locally closed-loop simple impedance control, and a completely passive, flexibly contained fluid volume to transmit impedance. It depends on designing the transmission to be suitably transparent, and the details of such a design are discussed below. A prototype is described in the following chapter.

6.2.2 High source impedance, dominant fluid dynamics

If the source actuator has high impedance, it may be advantageous to sculpt the transmission dynamics to improve interaction. The methods described in the previous

238

chapter could be used to design the details, but it is likely that compliance, damping and inertia all play an important role.

As in the previous design, feedback control can be done locally at the source actuator. This is consistent with the architectures explored in chapter 5. Impedance shapers are placed outside the control loop, between the force sensor and the environment. In this configuration, it is proposed that the transmission could also serve as an impedance shapaer. The dynamics of the fluid system are important in their direct physical effect on endpoint impedance as well as in their ability to permit more aggressive force feedback.

In general, it is expected from the previous chapter's analysis that impedance shaper compliance, damping, and inertia should be made nonzero to permit more aggressive force feedback. These are the parameters that are likely to be controlled to improve performance. Yet these are precisely the parameters to be minimized for a transparent design. Increasing these parameters is not likely to be difficult (though precisely implementing desired values might be). It is more challenging to minimize them.

Designing a fluid transmission with the goal of transparency identifies limiting dynamic parameters (minimum compliance, damping, inertia). With this information, additional dynamics can be more accurately implemented, and the transparent system can be augmented to shape the impedance. The rest of this chapter and the next focus mostly on designing a transparent transmission. Suggestions are offered for incorporating different dynamics as needed. Appendix B presents several proposed designs for an impedance shaping module that could be added to an otherwise transparent fluid transmission.

6.3 Design of mechanically transparent fluid transmissions

With transparency defined as a goal, it is helpful to focus on a specific configuration to understand the core design issues. A configuration is proposed for linear actuation, and a design model is derived. The next chapter describes the design, using this model, and testing of a specific prototype. A photograph of this prototype is shown in figure 6-1. It is used with a linear electromagnetic motor and affects linear forces and motion at the endpoint. Forces are transmitted through a fluid system that is contained in a flexible hose, permitting movement of the end-effector with a stationary source actuator. This chapter discusses some of the basic design decisions and governing calculations that led to the selection of this design configuration. Some fluid mechanics are reviewed in order to understand the tradeoffs concerning the geometry of this design, which are summarized at the end of this chapter. In the next chapter, the detailed design of the prototype in figure 6-1 is discussed.

6.3.1 Linear travel configuration

A linear actuator embodiment of the proposed actuation configuration is depicted schematically in Figure 6-2. The source actuator, responsible for generating force, motion, and impedance, is an electromagnetic linear motor. The stationary actuator's potentially large mass does not affect the output force density. The transmission is a completely passive and intrinsically low-impedance fluid system enclosed in flexible hoses. At the motor and at the endpoint, a piston-cylinder is used to transduce from the hydraulic to mechanical domain (and vice versa). The endpoint package, consisting only of the piston-cylinder and any necessary bearings and connected to the source actuator through flexible lines, is designed to have the force-to-mass advantages of hydraulics and the force-to-impedance properties of the electromagnetic actuator.

240
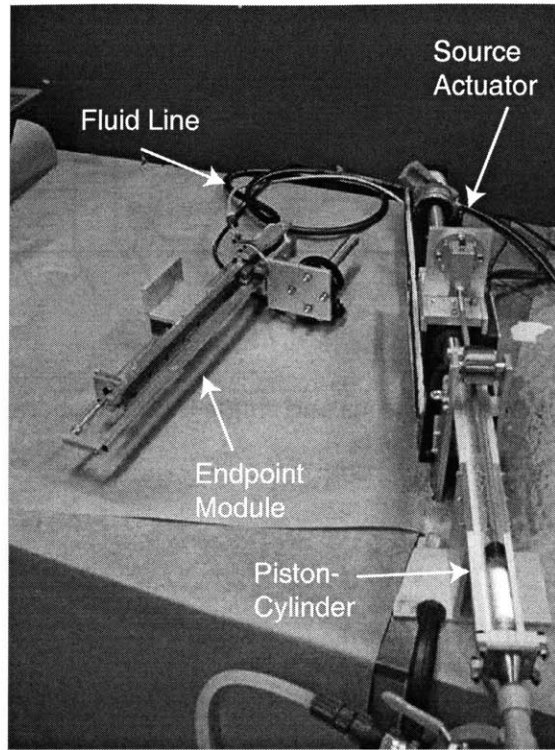
Figure 6-1: Photograph of the transmission system prototype developed in chapter7.
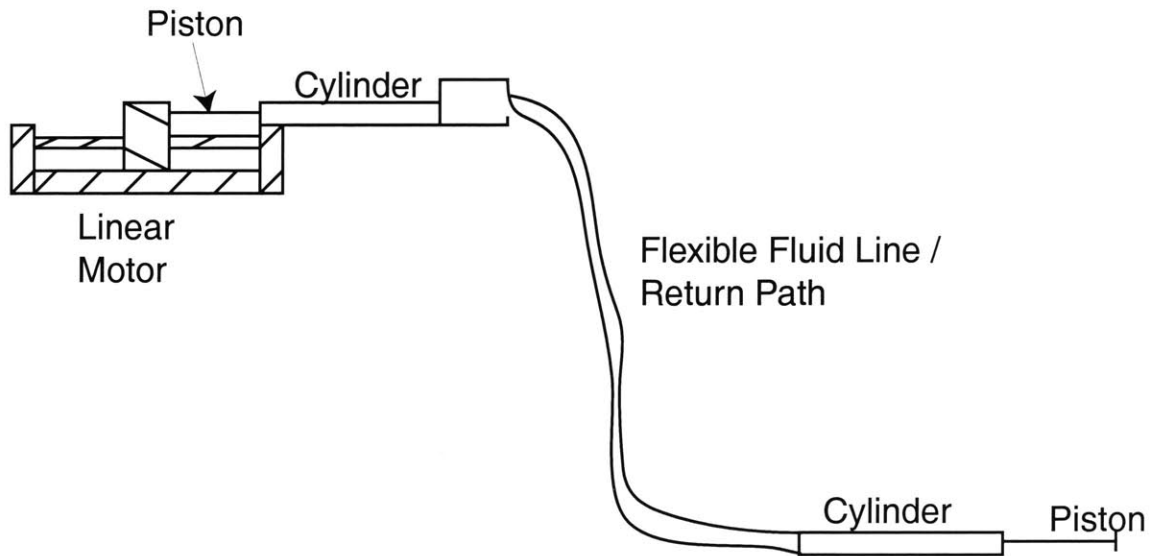
Figure 6-2: A rough schematic of a proposed linear actuation embodiment of a remote hydraulic transmission with an electromagnetic source actuator.
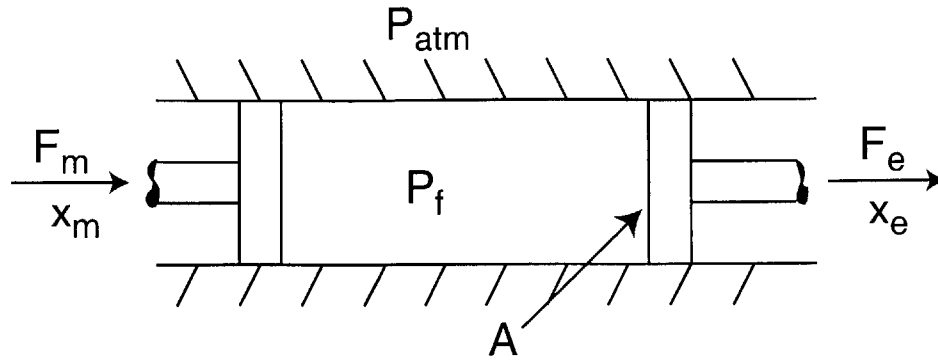
Figure 6-3: Simplified model of pistons and single fluid line, here modeled as a straight pipe.

Basic force transmission function

To illustrate the basic power transmission mechanism, figure 6-3 shows a simplified model of the proposed transmission system design. The transmission line is modeled as a straight pipe. On each end is a piston; one is connected to the motor, and the other to the environment. The motor force and position are represented by F_m and x_m, respectively, while the environment's force and position are represented by F_e and x_e, with all sign conventions positive in the direction of the arrows in Figure 6-3, as seen by the pistons. The cross-sectional area of both pistons and the pipe is A.

The symmetry of the transmission is apparent from Figure 6-3. This is a critical property; the transmission must transmit force and motion equally from motor to environment and from environment to motor. If x_e is fixed, and a positive force $F_m > 0$ is applied to the left piston, the pressure in the fluid is $\frac{F_m}{A}$ greater than atmospheric pressure (from a force balance on the left piston). The fluid is assumed incompressible, so x_m remains fixed. By a force balance on the right piston, $F_e = -F_m$. Thus the force is transmitted ideally from the motor to the environment. Here all dynamic effects, such as transmission delays, are neglected. If instead a negative force $F_e < 0$ is applied on the right piston, the identical result emerges.

242

6.3.2 Differences from servo-hydraulics

With the specific configuration of figure 6-2, the key difference between this design and conventional servo-hydraulic systems can be highlighted. In this system, pressure and energy come from an electromagnetic source that varies in real time with low endpoint impedance and that can equally handle energy flow in either direction. Alternatively, pressure and energy originate from the environment and flow to the source actuator, where they are stored or dissipated. The force or impedance signal command enters the hydraulic system from the same port as the energy from the actuator or environment; indeed the signal and energy flows are one and the same. This stands in stark contrast with the servo architecture described in chapter 2, which provides mechanical energy from a constant-pressure supply that is metered at (relatively) high impedance by an electromechanical servovalve. The source pump cannot accept backflow, and the servovalve must shunt any flow of energy into the hydraulic subsystem away to the drain line. Energy comes from the pressure source (the compressor) but the signal that commands output pressure and force comes from a separate source via the electromechanical servovalve, which has high intrinsic impedance.

Furthermore, the hydraulic portion of the proposed system and the mechanical-to-hydraulic (and hydraulic-to-mechanical) transduction are designed for low intrinsic impedance. No feedback is required around the hydraulic part of the system to reduce its intrinsic impedance, because it can be designed with physically low impedance.

6.3.3 Advantages

While the passive fluid transmission design was conceived principally to solve the force-mass-impedance problem as described above, this approach offers several additional characteristics that may increase its value for certain applications. Several are noted here.

Local feedback loop

Although the control scheme is somewhat independent of hardware design, the envisioned configuration has at least one advantage that might help to make control robust and effective. As mentioned above, the control loop may be able to be closed around only the local actuator's sensors, providing very high bandwidth. If the transmission is effectively made transparent, the interaction properties at the output should approach those at the actuator, so there is no need to close a control loop around the transmission's likely lower-frequency dynamics. This is an important consideration when using remote transmissions, because any system that transmits force over a large distance is likely to have non-negligible resonant dynamics, and be vulnerable to transmission delays. Closing a control loop around such dynamics can cause stability problems.

Intrinsic gearing

For clarity, the schematics above present motor- and environment-side pistons with equal effective area. If the areas are different, however, a gear ratio is provided between the two. This situation is depicted in Figure 6-4. As the volume of the incompressible fluid is assumed constant, the displacements of the two pistons must be related by the ratio of the two areas. The displacements x_m and x_e from some nominal position where both are zero are related by the following equation:

$$x_e A_e = x_m A_m \tag{6.1}$$

Similarly, the assumption that static pressure is constant in the system leads to the following equation:

$$\frac{F_m}{A_m} = \frac{F_e}{A_e} \tag{6.2}$$

When $A_m \neq A_e$, these two equations are characteristic of a gear reduction. By the usual convention that assumes gear ratios greater than unity when the motor torque
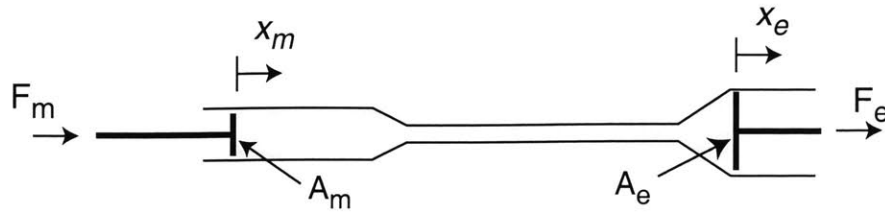
Figure 6-4: Different piston areas create a natural gear reduction across a fluid transmission, without extra hardware.

is amplified, the gear ratio is:

$$N = \frac{A_e}{A_m} \tag{6.3}$$

The implications of ideal gear reductions are addressed in section 2.2.2; however this configuration changes the calculus somewhat. With the source actuator stationary and somewhat remote, its weight is less relevant. Thus it is worth considering not only "gearing down" the motor in speed, but also possibly "gearing it up," in other words making $A_m > A_e$. There are compelling reasons not to do this: since actuator force is usually the limiting factor, the source motor generally has plenty of extra power, and the larger it is, the greater this excess. Larger actuators are usually more expensive and, obviously, take up more space. However, applications may exist that demand force and impedance performance that are better met by a "geared up" large actuator (with $N < 1$) than by a smaller motor with $N > 1$. Figure 2-4 suggests that there may be a benefit in reducing both reflected inertia and reflected viscous damping in using a larger source motor with smaller N. On the other hand, analysis in subsequent sections shows that the fluid transmission is likely to play a strong role in determining system damping and inertia, along with the source motor. If the motor's contribution to endpoint impedance is minimal, then there is no need to gear up a larger actuator. Other applications may find significant benefit in the cost and power savings of using a "geared down" motor (with $N > 1$), and may not suffer from the increased reflected actuator impedance properties this is likely to produce. This decision ultimately depends on the application; section 6.3.7 provides a summary of the design tradeoffs for use in various specific embodiments.

Actuation in hostile environments

Another potential advantage of the passive fluid transmissions proposed here is that, in their simplest form, they require very few and very simple components at the port of interaction with the environment. This is critical if the robot is expected to operate in certain hostile environments.

For example, functional MRI (Magnetic Resonance Imaging) is used to observe changes in a subject's brain activity as the subject moves a limb. The MRI protocol requires that the subject be inside a very strong magnetic field. Electromagnetic actuators and sensors with electronics are unable to operate properly under such conditions, and their presence can also obscure the image. Previously, imaging with active robot intervention has been limited to alternative scanning techniques such as Positron Emission Tomography (PET) [72]. When a passive fluid transmission is used, no sensors or actuators are required at the end-effector. The source actuator and sensor can be located a safe distance from the magnet, while the robot linkage itself can be made exclusively of materials that will not corrupt the image or malfunction inside the magnet. The fluid transmission requires only the lines, fittings, piston, cylinder, and seal, any bearings, and the fluid be inside the magnet. Each of these parts can likely be made of materials that are magnetically transparent, and construction of a prototype for these purposes is underway at the time of this writing.

Other hostile environments might also demand actuators that need no electronics, optics, magnetics, or electric power locally. A passive fluid transmission reduces the number of different types of parts that need be located on a robot linkage or at its end-effector. The parts that remain are simple mechanical and hydraulic fixtures. This characteristic could be useful in other, as-yet-unanticipated robot contexts.

6.3.4 Tensile force transmission

If a negative force $F_m < 0$ is applied on the left piston in figure 6-2, a force balance suggests that the internal pressure is $\frac{F_m}{A}$ less than atmospheric pressure. This suggests a problem, and intuition verifies it. Fluid cannot support force in tension; if tension

forces are applied, only atmospheric pressure can support these forces as the pressure in the closed fluid volume drops. This is undesirable for several reasons. If the pressure between the pistons is less than atmospheric pressure, then any leakage across the seals will put air into the system, making the fluid compressible. Similarly, if the local pressure anywhere in the fluid drops below the saturation pressure for its temperature, it evaporates locally and cavitation results [44]. Even in the absence of these problems, the tension force the fluid supports is limited absolutely by the product of atmospheric pressure and the piston area. Unless the system is designed for only extremely low operating pressures (atmospheric pressure is approximately 15 psi), this force is dramatically less than the compressive force the system can bear.

Applying substantial tension forces to the fluid chamber is, without doubt, a bad idea. A versatile transmission must, however, be able to transmit force in both directions. Two solutions are proposed below.

Opposing fluid paths

The system in Figure 6-3 is unable to sustain tension forces because when the left piston is pulled to the left, there is nothing there but air to support it. This can be changed if a second fluid chamber is added. Figure 6-5 shows a simplified version of the system with dual fluid paths. When x_e is fixed and a force $F_m > 0$ is applied to the left piston, chamber A is active and its pressure rises. Conversely, when a force $F_m < 0$ is applied, chamber B supports the pressure. By the symmetry in the system, if forces F_e are applied, they are similarly transmitted to the motor.

One disadvantage of this system is that additional seals are required; in addition to the two piston seals, each of the piston rods must also be sealed. Leakage across the piston seals changes the relative volume between chambers A and B, and therefore the relationship between x_m and x_e is changed. Leakage out the piston rod seals actually exits the system, and something must take its place. If there is no fluid makeup system, this is likely to be air. Fluid makeup is discussed in the following chapter. Another disadvantage is that this system requires roughly twice as much fluid, increasing the endpoint mass of the system.
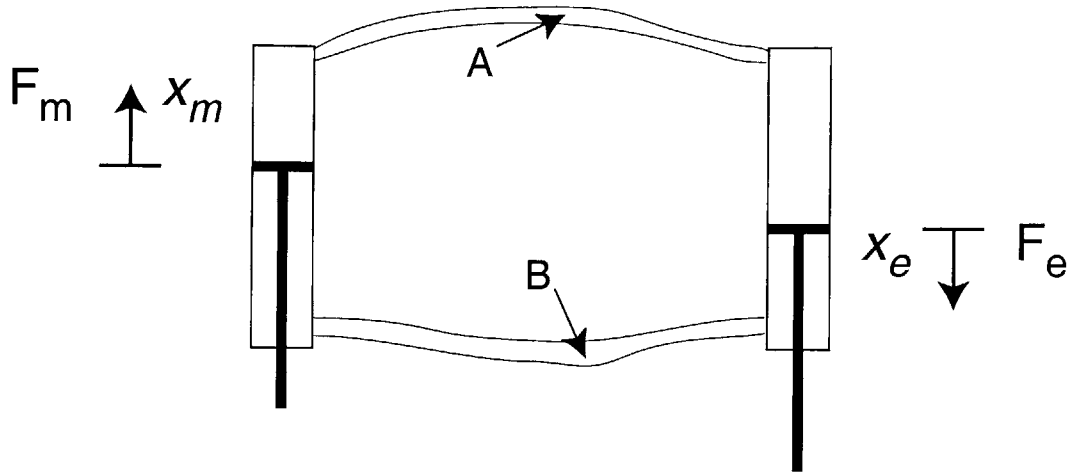
Figure 6-5: Passive hydraulic system with dual fluid paths, for bilateral force transmission.

One possible advantage of this configuration is that the system can be elevated to a bias pressure without producing any net force output. Only the differential pressure affects piston force, not the absolute pressure. This can be used to compress any remaining air and reduce its impact on stiffness. Elevating to a bias pressure can also prove useful for fluid makeup.

If this configuration is used, and the two piston-cylinders are not identical, care must be taken in sizing the piston rods to ensure that the total volume of chamber B remains constant as the pistons move.

External bias pressure

Forces can be supported in both directions with a single fluid chamber if a neutral bias pressure is mechanically introduced to the system, providing a margin for the fluid pressure to safely drop from neutral while remaining above atmospheric pressure. This is shown schematically in Figure 6-6, where the force that introduces the bias is denoted by F_{bias}. The forces applied at the two pistons must be equal in magnitude, so that their influence on the total system cancels out and the system is stationary if $F_m = F_e = 0$. A force balance on either individual piston, however, reveals that the interior system pressure must compensate for the bias force as well as any applied
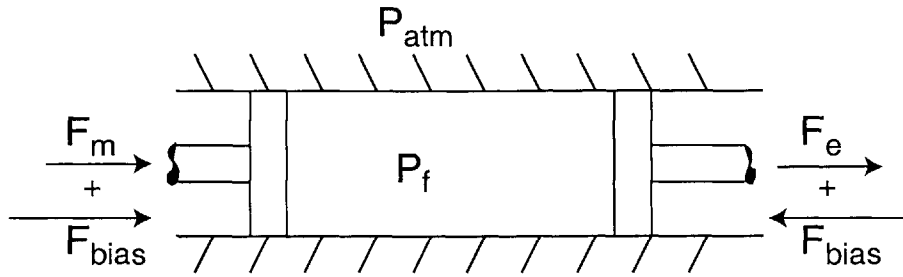
Figure 6-6: Simple fluid line model with equal bias force applied at both pistons.
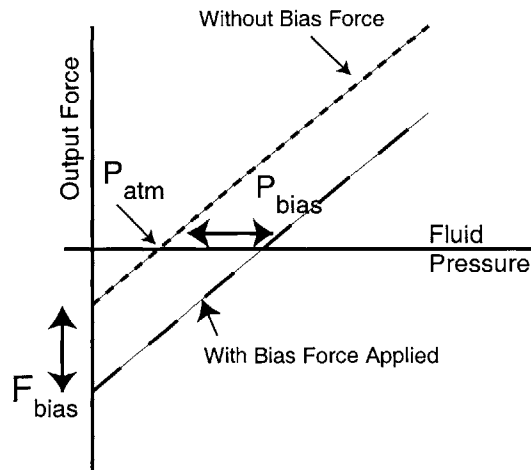
Figure 6-7: The effect of adding a bias force, as depicted in figure 6-6, to a single-path hydraulic transmission.

force (F_m or F_e).

If F_{bias} is constant, the result is a uniform increase in internal pressure, equal to $P_{bias} = \frac{F_{bias}}{A}$. Figure 6-7 illustrates the effect of a bias force on the static pressure versus output force. With no bias force, the output force is zero when the fluid is at atmospheric pressure P_{atm}. With the bias force, the pressure is offset at all frequencies by P_{bias} so that the output force is zero at $P_{atm} + P_{bias}$, and the system can support F_{bias} in tension without fluid pressure dropping below P_{atm}. The higher the bias force, the higher the tension force that the system can support.

The most obvious disadvantage of this approach is that it increases the ratio

of maximum pressure to maximum output force. This is evident from Figure 6-7. Components are often limited by some maximum permissible pressure P_{lim}, and this approach requires higher pressures for the same output force than the dual fluid chamber approach described in Section 6.3.4. If the actuator is to be capable of equal forces in both directions (this is not necessarily the case), and the pressure is to stay above P_{atm}, then the maximum bidirectional output force capacity is achieved when the bias force is such that P_{bias} is exactly half of the pressure limit. In this case, the maximum force in each direction is $\frac{P_{lim}}{2}A$, which is equal to $P_{bias}A$. This is half of the unidirectional force limit. When additional compressive force is applied ($F_m > 0$), at maximum force the pressure is P_{lim}. When force is applied in the tensile direction, at maximum force the pressure is equal to the atmospheric pressure.

This approach requires only a single chamber of fluid, and therefore roughly half the fluid mass of the alternative presented above. However, it requires that a force be applied continuously at each piston. Analysis and control are more straightforward if this force is a constant. On one side, the actuator can be used directly to apply the force if additional authority is available. This solution is unavailable on the environment side. Some alternative means must be used, depending on the configuration of the robot, but any such approach is likely to add both mass and complexity. Several possibilities for a particular application are discussed in the following chapter.

An advantage of this approach is that there are only two moving seals, one at each piston. If fluid leaks out of the system without being replaced, the volume of the system is reduced, and in theory, no air enters the closed system. The relationship between x_m and x_e may change, but large reductions in stiffness are avoided. Precise positioning accuracy may be sacrificed, but the force and impedance transmission should be roughly constant. This makes the system more robust to leakage than the dual-chamber solution. If the volume of the fluid must stay constant, a fluid makeup system can be employed to replace any fluid that is lost to leakage. Fluid makeup is discussed in the following chapter.

6.3.5 Fluid dynamics

In the initial descriptions above, uniform pressure throughout the fluid is assumed, and force relationships between the pistons are considered in the context of static pressure. Of course an actuator requires movement as well as force transmission, and the movement of fluid produces its own dynamics that affect endpoint impedance and efficiency. This section analyzes the effects of fluid dynamics and leads to an explanation of the tradeoffs between geometric and fluid parameters and the transparency of the transmission system.

When a fluid flows along a surface, friction produces a velocity-dependent pressure drop. This is analogous to mechanical friction in its velocity dependence, but might have different characteristics. Similarly, when the fluid is accelerated a pressure change results from the inertia of the fluid. The geometry of the system provides the basis for model predictions of friction and inertia.

Energy equation for pipe flow

Analysis of damping in the fluid lines begins with an analysis of head loss. Steady, incompressible, frictionless flow between any two points in a pipe is governed by the Bernoulli equation:

$$(\frac{P_1}{\rho} + \frac{v_1{}^2}{2} + gz_1) - (\frac{P_2}{\rho} + \frac{v_2{}^2}{2} + gz_2) = 0 \tag{6.4}$$

The terms in the first set of parentheses concern the first point, and the terms in the second set concern the second point. P_1 and P_2 are pressures, v_1 and v_2 are bulk fluid velocities, ρ is the fluid density, g is the gravitational constant, and z_1 and z_2 each represent elevation (as measured from the same baseline). This is perhaps the most familiar form for analysis of internal flow.

A similar form that includes the effects of fluid viscosity can be derived from the energy equation [44]:

$$(\frac{P_1}{\rho} + \alpha_1 \frac{v_1{}^2}{2} + gz_1) - (\frac{P_2}{\rho} + \alpha_2 \frac{v_2{}^2}{2} + gz_2) = h_{lt} \tag{6.5}$$

This is the basic equation used for analysis of pressure, flow, and gravity for this system. α_1 and α_2 are kinetic energy coefficients that depend on the area integral of the flow velocity profile. α can range from 1.0 to 2.0 depending on the flow profile and Reynolds number. At high Reynolds number, α approaches 1.0. For laminar pipe flow, $\alpha = 2.0$. Usually the kinetic energy change is most relevant for developing flows at moderate Reynolds number [44].

Each collection of terms in parentheses in equation 6.5 represents the mechanical energy per unit mass at a particular cross section. The total head loss h_{lt} expresses the energy lost to friction between these two cross-sections, and has units of energy per unit mass. For this analysis, points 1 and 2 each correspond to one of the pistons, or more accurately to the fluid immediately adjacent to the pistons. This is because the fluid dynamics only matter in this application as they appear at the pistons, where the fluid system interfaces to the mechanical system via either the actuator or environment. The velocities v_1 and v_2 are the piston velocities, because the velocity of each piston is equal to the average velocity of the fluid immediately adjacent to it. Similarly, the pressure at each piston is proportional to the force applied to the piston:

$$P_i = \frac{F_i}{A_i} \tag{6.6}$$

for $i = 1$ or 2 (where P_i is referenced to atmospheric pressure). This is determined from a force balance.

Various sources of frictional losses under various conditions produce different expressions for head loss, and each of these terms can be included in h_{lt}. If the head loss depends linearly on velocity, this loss is in the form of linear viscous damping. In many cases it is nonlinear and depends on higher powers of velocity. Details of fluid damping for several relevant system features are provided below and in the following chapter.

Figure 6-8 shows two pistons connected by a pipe of different diameter. This is a simple fluid dynamics model for this problem. The left piston, piston 1, is associated with P_1, v_1, α_1, and z_1, and has area A_1. The right piston is associated with P_2, v_2,
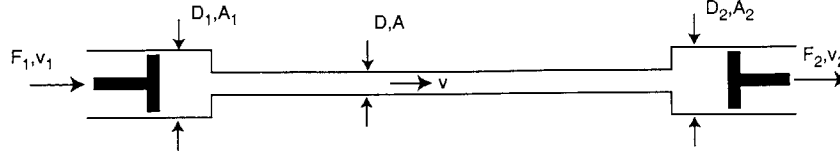
Figure 6-8: Model of fluid system with two pistons and a pipe of different diameter.

α_2, and z_2, and has area A_2.

The general form of equation 6.5 can often be simplified for particular examples of the fluid transmission. To simplify damping calculations, it is convenient to consider force F_1 applied to produce movement at constant velocity at piston 1, with no resisting forced applied at piston 2 ($F_2 = 0$). From equation 6.6, $P_2 = 0$. If both pistons are at the same elevation, gravity has no effect, and the terms gz_1 and gz_2 can be eliminated. Because the fluid is assumed incompressible, if both pistons have the same cross-sectional area, conservation of mass dictates that both pistons travel at the same velocity, so $v_1 = v_2$. Immediately adjacent to the piston, fluid is forced to follow a rectangular flow profile with each streamline traveling at the same velocity as the piston. Thus when the areas are equal, the flow profiles and Reynolds numbers are identical, and $\alpha_1 = \alpha_2$, eliminating the fluid dynamics terms. Under these assumptions ($F_2 = 0$, $z_1 = z_2$, $A_1 = A_2$), equation 6.5 is simplified to the following:

$$P_1 = \rho h_{lt} \tag{6.7}$$

Under these conditions, the entire pressure drop produced by applying force $F_1 = P_1 A_1$ is taken up by friction, expressed by the head loss.

When the assumptions in the previous paragraph do not hold, the relative magnitudes of the terms in 6.5 can be compared to gauge the significance of each effect. If A_1 exceeds A_2, and therefore v_2 exceeds v_1, the impact of the fluid dynamic terms can be checked by comparing:

$$|\alpha_1 \frac{v_1^2}{2} - \alpha_2 \frac{v_2^2}{2}| \tag{6.8}$$

with the magnitude of the major head loss, at several velocities within the expected range. For turbulent flows, α can be determined based on Reynolds number and an empirically determined flow profile (see section 8-5 in [44]). Alternatively, the size of this term can be bounded by recognizing that α is between 1.0 and 2.0. Similarly, computing the maximum magnitude of the gravity terms is straightforward when the maximum expected difference in elevation is assumed for $(z_1 - z_2)$.

Damping, linear and nonlinear

Friction produces the head loss h_{lt} in equation 6.5. For systems with internal flows, head loss is a combination of major losses h_l, due to damping from the straight lengths of pipe, and minor losses h_{lm}, due to obstructions in flow such as changes in area, corners and bends, and orifices:

$$h_{lt} = h_l + h_{lm} \tag{6.9}$$

The labels "major" and "minor" do not imply a difference in magnitude; friction can be dominated by either major or minor losses, depending on the system. The two types of losses are treated separately because major losses from fully developed flow in straight pipes can be computed in closed-form, while minor losses are determined from empirical data. Minor loss may be expressed as:

$$h_{lm} = K \frac{v^2}{2} \tag{6.10}$$

where v is the average velocity and K must be determined experimentally. Alternatively, head loss is sometimes expressed as:

$$h_{lm} = f \frac{L_e}{D} \frac{v^2}{2} \tag{6.11}$$

where L_e is an effective length of pipe, D is the pipe diameter, and f is a *friction factor*. Experimental results found in any fluid mechanics textbook provide estimates of minor head loss for most common geometries (for example see section 8.7 in [44]).

Once the head loss has been estimated and used in equation 6.5 (or equation 6.7 if the necessary assumptions hold), this expression can be rearranged in order to compute the effective damping of the transmission, as seen at either piston. If the piston areas are the same, the effective damping for forces applied at either piston is identical. If the areas differ, the apparent damping at the two pistons differs by a function of the ratio of the areas. The apparent friction at a piston is a relationship between the force and velocity at that piston. Thus in equation 6.5, each fluid velocity must be converted into the appropriate piston velocity, and the pressure drop must be related to piston force. As an example, the effective damping at piston 1 due only to major losses in the intermediate pipe in Figure 6-8 is computed for laminar flow. Friction in the entrance and exit to the pipes, as well as friction in the pistons themselves, is neglected. The major head loss, as a function of local average velocity v, local pipe area D^2, and pipe length L, fluid density ρ and viscosity μ is (from equation 8.31 in [44]):

$$h_l = 32\frac{L}{D^2}\frac{\mu}{\rho}v \tag{6.12}$$

Because the volume flow rate in the pipe is assumed the same as in the piston, the velocity can be related to the piston velocity v_1 via the following expression:

$$v_1 A_1 = vA \tag{6.13}$$

If dynamic flow effects and gravity are neglected, and piston 2 is assumed massless and frictionless and allowed to move freely such that $P_2 = 0$, then by equations 6.7, 6.12, and 6.13:

$$P_1 = 32\frac{L}{D^2}\mu\frac{A_1}{A}v_1 = 32\frac{L}{D^2}\mu\frac{D_1{}^2}{D^2}v_1 \tag{6.14}$$

By observing that $F_1 = P_1 A_1$, this can be converted to:

$$F_1 = B_{eq}v_1 \tag{6.15}$$

255

The friction is in the form of linear damping, with the equivalent damping coefficient:

$$B_{eq} = 8\pi\mu L(\frac{D_1}{D})^4 \tag{6.16}$$

The major damping at the piston is proportional to the *fourth power* of the ratio of piston diameter to pipe diameter (or the square of the area ratio $\frac{A_1}{A}$) in the laminar region. In other words, making the pipe smaller increases damping very dramatically in the laminar region.

If the flow in the pipe is turbulent, the major loss must be computed with a different expression. In general, transition from laminar to turbulent flow is expected with Reynolds number around 2300. The Reynolds number is defined as:

$$Re = \frac{\rho v D}{\mu} \tag{6.17}$$

When flow is turbulent, equation 6.12 does not apply, and an alternative expression for h_l is required. Major head losses can be defined as:

$$h_l = f\frac{L}{D}\frac{v^2}{2} \tag{6.18}$$

The friction factor f must be determined from a Moody diagram and depends on the Reynolds number as well as the relative roughness (commonly denoted $\frac{e}{D}$) of the pipe material. Rubber and plastic pipes, as used in this application, are quite smooth such that $\frac{e}{D}$ approaches zero. To proceed symbolically rather than rely on a Moody diagram, it is helpful to use an approximation for smooth pipes. For Reynolds numbers between 4,000 and 100,000 the Blasius equation can be used [44]:

$$f = \frac{0.3164}{Re^{0.25}} \tag{6.19}$$

Substituting equations 6.17 and 6.19 into 6.18 produces:

$$h_l = \frac{0.3164\mu^{0.25}}{\rho^{0.25}}\frac{L}{D^{1.25}}\frac{v^{1.75}}{2} \tag{6.20}$$

Applying equations 6.6 and 6.13, the force-velocity relationship can be derived:

$$F_1 = \frac{\pi}{8} 0.3164 \mu^{0.25} \rho^{0.75} L \frac{D_1^{5.5}}{D^{4.75}} v_1^{1.75} \tag{6.21}$$

The damping force depends on the piston velocity to the 1.75 power. There is an even stronger dependence on both the piston (D_1) and line (D) diameters as in the laminar case.

Although turbulent flow produces nonlinear damping, this may not be problematic in design or control of human-interactive systems. The nonlinearities are "soft" (they lack discontinuities) and behavior is smooth. Well-behaved viscous friction can even help to suppress unwanted high-frequency vibration. This contrasts with problematic mechanical friction, which is characterized by discontinuity and even local negative damping [5, 6, 15].

This process can be repeated to determine the net damping effect of each particular source of head loss, or all head losses can be included simultaneously and a single equivalent damping can be computed. In the next chapter these calculations are used in the prototype design.

When damping losses are computed, the pipe is assumed perfectly rigid and stationary. This is not strictly true, but the assumption should provide a good approximation provided the compliance is not excessively large relative to the compliance of the source actuator, producing a large difference in velocities at the two pistons. For most operating conditions of this prototype, this assumption should hold. Potential changes to the damping as the system is moved (i.e. the hose is bent) are also ignored. The implications of this second assumption are less obvious, and depend greatly on the the particular embodiment of the system. If all design calculations are based on a straight pipe, but in the actual design the hose must wrap tightly around corners, the actual damping may be much higher than estimated. The best way to deal with this is to make more than one estimate based on the anticipated usage of the system, in order to establish reasonable upper and lower limits on the friction estimates.

Significantly, equation 6.5 provides an expression for the relative pressure drop

between two points as fluid flows, here assumed to be the two pistons. The effective damping does not depend on the absolute pressure. This is critical for analysis of this application because unlike many fluids problems, the pressure does not itself drive the flow. Pressure and flow rate are both determined by the inputs at both sides of the line. Thus pressure and flow rate vary independently, and any estimate based on flow velocity must consider variations in absolute system pressure due to forces applied at the pistons.

Fluid inertia

Like any substance with mass, fluid exhibits inertial properties when accelerated. In a fixed-volume system like this one, the fluid density is one key factor in determining inertia. However, the geometry of the system also plays a critical role in determining endpoint inertia.

For the system in figure 6-8, the equivalent fluid mass as seen at piston 1 (or piston 2, because the system is symmetric) can be determined by an energy calculation. Assuming the fluid is incompressible, at any instant in time, the total kinetic energy of the fluid can be expressed as the energy of an equivalent mass m_{eq} moving at the piston velocity v_1:

$$m_{eq}\frac{v_1{}^2}{2} = m_{cyl}\frac{v_1{}^2}{2} + m_{line}\frac{v^2}{2} \tag{6.22}$$

m_{cyl} is the mass of the fluid in the two cylinders. If the cylinders have equal size, this quantity is constant even as the pistons move. m_{line} is the mass of the fluid in the line, which is also a constant quantity. This fluid moves with velocity v, which is related to the piston velocity by equation 6.13. Combining equations 6.13 and equation 6.22 produces:

$$m_{eq} = m_{cyl} + m_{line}(\frac{A_1}{A})^2 \tag{6.23}$$

The total fluid mass in the cylinders is simply the volume (of one cylinder) times the density:

$$m_{cyl} = \rho L_1 \frac{\pi D_1{}^2}{4} \tag{6.24}$$

258

where L_1 is the length of a cylinder. Similarly:

$$m_{line} = \rho L \frac{\pi D^2}{4} \tag{6.25}$$

Substituting into equation 6.24 and putting the cross-sectional areas in terms of diameter produces:

$$m_{eq} = \rho L_1 \frac{\pi D_1{}^2}{4} + \rho L \frac{\pi D_1{}^4}{4D^2} \tag{6.26}$$

At least in the prototype designed here, the fluid inertia due to the line dominates. The line is almost 10 times as long as the cylinders, and the diameter D is considerably smaller as well. The line inertia is inversely proportional to the square of the line diameter. So when the line is made smaller in diameter, the apparent inertia increases, even though there is less fluid. Conversely, when the line is made larger, the inertia decreases.

Effect of damping and inertia on endpoint impedance

The friction and apparent inertia determined by the calculations above are bulk properties of the fluid transmission, and under the assumption of infinite stiffness for the transmission, contribute directly to the endpoint impedance of the system. Figure 6-9 helps to illustrate this point. If the uncontrolled actuator is modeled as a mass m_{act} retarded by some linear or nonlinear friction function $f_{act}(v)$, and the transmission is modeled as a mass m_f retarded by some friction function $f_f(v)$, and the actuator and transmission are joined so as to have the same operating velocity v, then the two inertias and the two friction functions add directly. This is shown by the bond graph in Figure 6-10. The total inertia seen by the environment is the sum of the actuator and transmission inertias:

$$m_{port} = m_{act} + m_f \tag{6.27}$$

Similarly, the total friction seen at the port is the sum of the individual friction functions:

$$f_{port} = f_{act} + f_f \tag{6.28}$$
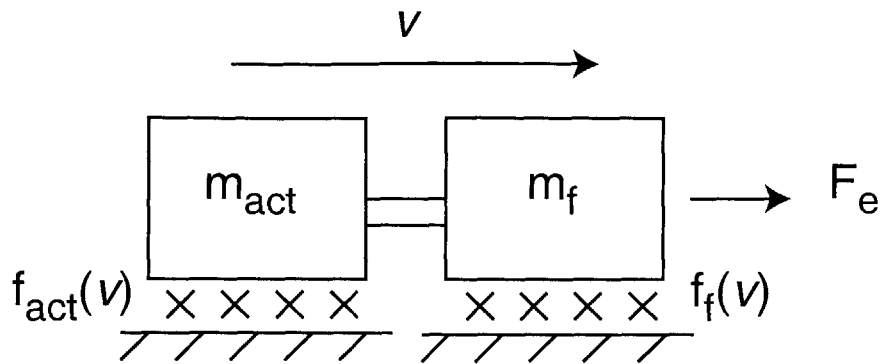
Figure 6-9: Assuming zero compliance, the actuator and fluid mass move as one. Thus the endpoint impedance includes friction and inertia due to both the fluid and the actuator.
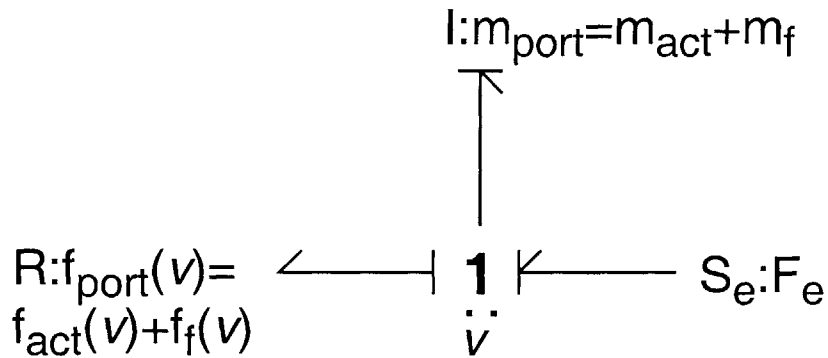
Figure 6-10: Assuming zero fluid compliance, the actuator and fluid pistons move with the same characteristic velocity v, so that the actuator and fluid friction contribute to the endpoint impedance.

This provides further support for the statement that minimizing transmission inertia and friction is desired. The ungeared electromagnetic actuator is chosen specifically for its low inertia and friction; to be effective, the transmission must preserve these qualities.

Leakage

Any system with differential pressure across sliding seals is vulnerable to leakage. This is a particular concern in this application, as low friction is desired at the seal, in order to preserve low impedance performance. In general, more effective seals

260
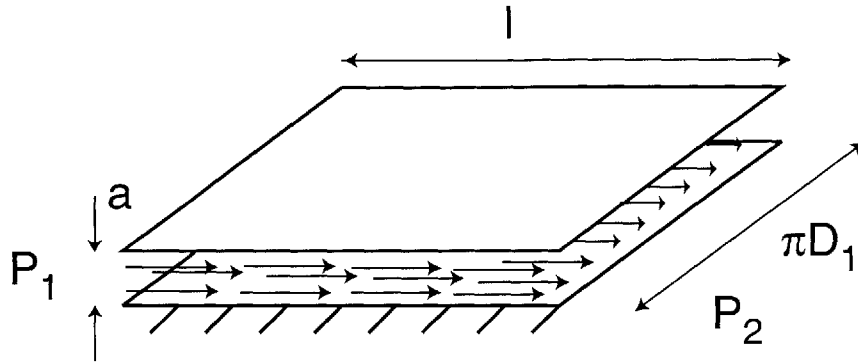
Figure 6-11: Model of leakage flow across a seal.

have higher friction. A discussion of various seal types and their friction properties is provided in the next chapter.

Assuming there is a small fixed gap between the seal and the cylinder, and a differential pressure across the seal, the leakage flow can be treated as forced internal flow. Because the gap is small relative to the piston diameter, this can be modeled as flow between two parallel plates, depicted in figure 6-11. The flow depends on the duct geometry, the fluid viscosity, and the differential pressure between the two sides. Section 8-2 in [44] provides the equation for the volumetric leak rate Q:

$$Q = \frac{\pi D_1 a^3 \Delta P}{12 \mu l} \tag{6.29}$$

D is the piston diameter, a is the dimension of the gap, and l is the length of the piston. The pressure differential is:

$$\Delta P = P_1 - P_2 \tag{6.30}$$

This rate must be doubled to compute the total leakage out of the two pistons. The gap is the most important dimension in controlling leakage flow, as the volumetric flow rate is proportional to the third power of a. If the gap size is pressure-dependent, as it may be for some elastomeric seals, the calculation is more complicated.

261

6.3.6 Line compliance

All analysis to this point has assumed that the fluid is perfectly incompressible, and the system perfectly rigid, resulting in identical bulk flow rates at both pistons. For the purposes of this design, the first assumption is probably sound, provided that air is almost completely eliminated from the system. The second assumption, however, is more tenuous.

It is important to draw a distinction between the stiffness (or flexibility) of the hose in bending, and the stiffness (or compliance) in response to changing fluid pressure. Ideally, the hose should be infinitely flexible and have zero stiffness in bending, such that the hose itself does not introduce any forces in response to movement of the end-effector with respect to ground. On the other hand, the hose should ideally be infinitely stiff radially, so as not to expand when the fluid pressure increases. This is necessary to fulfill the requirement of transparency, as defined at the beginning of this chapter. Unfortunately, the radial rigidity and bending stiffness are not generally independent; indeed if the hose were isotropic, these objectives would be directly contrary. In fact, hoses designed for all but the lowest pressures have radial reinforcement to improve both radial stiffness and pressure capacity. Nevertheless, many hoses that are radially stiff are also resistant to bending, and decisions on this tradeoff must be made with specific hose types in mind.

Because hoses can be reinforced in so many different ways, many of them involving braided fiber, it is difficult to develop an accurate general model for line compliance. However, approximating the line as a thin-walled cylindrical pressure vessel can provide insight into the effects of changing the system's geometry. Returning to the notation of figure 6-8, the pistons are assumed radially rigid and the pipe wall is assumed to have thickness t. Piston 2 is assumed fixed, and the pipe is assumed to be free to expand and contract such that it does not support longitudinal stress. The stiffness at piston 1 can then be predicted. The hoop stress σ that develops in the pipe in response to an applied pressure P is given by:

$$\sigma = \frac{PD}{2t} \tag{6.31}$$

If the pipe is assumed to be made of a linear isotropic material (although a reinforced hose almost certainly is not) with modulus of elasticity E, and the radial compression of the pipe wall is neglected, this is equivalent to:

$$\sigma = E\frac{\Delta D}{D} \tag{6.32}$$

Combining these two equations yields:

$$P = 2tE\frac{\Delta D}{D^2} \tag{6.33}$$

By equating the change in fluid volume at the piston, $A_1\Delta x_1$ and in the pipe, ΔAL, the following relationship is determined:

$$\Delta D = \sqrt{\frac{D_1{}^2\Delta x_1}{L}} \tag{6.34}$$

Assuming that the pressure comes from an applied force F_1:

$$P = \frac{4F_1}{\pi D_1{}^2} \tag{6.35}$$

Combining the previous three equations and rearranging produces:

$$F_1 = \pi tE\frac{D_1{}^3}{D^2}\frac{1}{2\sqrt{L}}\sqrt{\Delta x_1} \tag{6.36}$$

Equation 6.36 provides several useful pieces of information. The apparent stiffness decreases if the pipe diameter is increased, even if the pipe thickness is increased proportionally, because the dependence on the diameter is stronger than on the thickness. The stiffness also decreases if the line length is increased. Both of these facts are explained by the fact that these changes increase the line volume, and thus comparable changes in percent volume lead to larger changes in total volume. The apparent stiffness increases with the third power of the piston diameter. The stiffness is nonlinear (and "softens") in this model, as force increases with the square root of displacement.

The model softens further when it is noted that increasing pressure likely decreases the thickness t through radial compression.

However, it is unlikely that the assumptions required to use the thin-walled equations will apply in all cases related to this transmission design. As noted above, most of the hose types considered (see the next chapter) are anisotropic. Furthermore, many have walls that are too thick relative to their diameter to permit this approximation.

Compared to the damping and inertia, the line compliance is difficult to estimate. Design can be guided by the knowledge that as less fluid is displaced with changes in pressure, the higher the line stiffness. All else being equal, compliance increases with line length and line diameter, and decreases with line thickness, material stiffness, and piston diameter. No serious attempt is made to accurately predict the line compliance of the prototype. Instead, this is measured. In some cases, specifications from the hose manufacturer might also prove useful.

6.3.7 Design tradeoffs

When viewed in aggregate, the fluid transmission model described above reveals some important tradeoffs that may not be obvious to engineering designers. The relevant relationships are summarized here, first in terms of device-level performance, and then in terms of several critical geometric parameters. Both sections provide similar information from different perspectives, each of which is useful for design.

Leakage versus port impedance

The leakage rate, predicted by equation 6.29, is inversely related to the port impedance for several reasons. In other words, changing the system to decrease the leak rate often increases the port impedance.

The most effective way to reduce leakage (as shown by equation 6.29) is to reduce the effective gap size a. However, this almost always increases the seal friction, which directly contributes to endpoint friction (this is discussed for specific seal designs in

the next chapter).

An alternative way to decrease leakage is to increase the fluid viscosity μ. However, this also produces a proportional increase in the linear damping coefficient, given by equation 6.16, as well as a weaker increase in friction in the turbulent region, shown by equation 6.21.

A third way to decrease leakage is to increase the effective piston or seal length, given by l in equation 6.29. This reduces the ratio of travel to total cylinder length and increases the size and mass of the endpoint package.

A final way to decrease leakage is to change the piston diameter D_1. Because the pressure differential ΔP is inversely proportional to the piston area (and the square of the diameter D_1), the leak rate is inversely proportional to D_1. Increasing D_1 reduces the operating pressure and therefore the leak rate. If D_1 is increased without a commensurate increase in D, however, the effective damping (equations 6.16 and 6.21) increases with at least the fourth power of D_1, and the equivalent inertia (equation 6.26) increases with the square of D_1. These problems can be mostly rectified by increasing the line diameter D to preserve a constant ratio $\frac{D_1}{D}$, but this increases the line compliance as well as the total mass and weight of the fluid lines. The effects of changing D and D_1 are summarized in the next section.

In summary, reducing the leakage by changing the system geometries or the fluid properties has an adverse effect on the port impedance through some combination of the following changes: increased seal friction (via a reduction in seal gap), increased damping coefficient (via an increase in viscosity), increased cylinder size (via an increase in piston length), increased reflected line damping and inertia (via an increase in piston diameter), and increased fluid line compliance and weight (via an increase in both piston and line diameters). In a practical design, each characteristic must be traded off against the others.

Changing D and D_1

The damping and inertia depend more strongly on the line and piston diameters than on anything else. Because of the importance of these parameters, in this section the

practical effects of modifying each independently are analyzed.

If the piston diameter D_1 is increased, the operating pressure to meet force requirements is decreased. Thus although the cylinder must be larger, its pressure loading is reduced. It is unclear whether this is likely to lead to a net increase or decrease in piston weight. By reducing the operating pressure, the leakage is decreased proportionally, and the line stiffness is increased (increasing D_1 relative to D is an effective gear reduction to the properties of the line). However, from equations 6.16, 6.21, and 6.26, the damping and inertia are increased.

If the line diameter D is increased, the line inertia and laminar and turbulent damping are all decreased with strong dependence, as shown by equations 6.16, 6.21, and 6.26. The leakage rate is unchanged. The total system weight and size are increased. If the hose type is not changed, a larger hose is almost certainly less flexible.

These tradeoffs summarize the application of simple fluid mechanics to the configuration in question. In the next chapter, they are used to design the details of the prototype in figure 6-1. Characterization experiments are then described that evaluate the accuracy of these models in search of any important physics that have not yet been considered.

Chapter 7

Fluid Transmission Design Example and Validation

While the previous chapter provides a model and governing equations for the design of an intrinsically low-impedance hydraulic transmission, this chapter presents the practical details related to the design of a prototype, as well as test results to check the model. Specific components are discussed in the context of their influence on system performance. The prototype was designed to provide a vertical therapy robot actuator, as described in chapter 1.

7.1 Specifications and goals

The specifications for the vertical actuator are provided in chapter 1, and repeated here for convenience. The force capacity should be at least 65 N in one direction and 45 N in the other. The endpoint inertia and total mass should each be limited to 2 kg. The static friction should ideally be less than 2 N, and the device should be able to render apparent stiffness between 0 and 2 kN/m. Chapter 2 describes two prototypes, a direct-drive linear motor and a screw-driven module. The first achieves the force and impedance specifications (with the exception of slightly more Coulomb friction than desired), but is more than four times too massive. The second is also almost four times too massive and also has high impedance, but with higher force capacity.

The control strategies developed in chapter 4 bring its impedance to specification, but its mass is still too large.

The objective here is to apply the design detailed in chapter 6 to create an actuator that can satisfy the force and impedance requirements with much lower endpoint mass. As this is the first embodiment of a new design, there are a number of specific issues to explore. A central goal is to determine whether the model in the previous chapter is reasonably complete, or whether any important part of the governing physics has been neglected.

Specifically, a key objective for the prototype was to verify the capability of the design to transmit power and impedance through a flexible line. A flexible hose length of 3 m was specified to provide sufficient distance from the robot end-effector to the source actuator. The flexibility of the hose is an important feature that must be traded off against other properties such as hose size and radial rigidity. As a first comparison, properties similar to those of the electrical cables that power comparably-sized electromechanical actuators were targeted. For actuators that produce appropriate force levels for this application, these cables typically have outer diameters on the order of 1 cm, and minimum bend radii of approximately 1 to 2 in (2.5 to 5 cm). In a serial robot that uses electromechanical actuators, such cables are unavoidable and are not generally problematic. Thus if the hydraulic lines can be designed to have similar size and flexibility, they should not introduce significant undesired behavior.

Beyond the verification of the expected force transmission capabilities, perhaps the most important role of the first prototype was to provide a platform to test the "endpoint feel" of the actuator/hydraulic transmission system. The feel, quantified throughout this thesis in terms of the mechanical impedance, is difficult to predict with certainty when a new design is created. Hard nonlinearities and other unmodeled phenomena can have a dramatic impact on the endpoint impedance and cause it to differ substantially from predictions. In other words, the designer of a human-interactive system can never be absolutely certain that it will feel right until he builds a prototype and feels it in action. Particular attention is paid to the friction and leakage properties of the seals and the working fluid, which are expected to be

the critical characteristics in achieving an effective low-impedance transmission.

7.2 Prototype design

The basic prototype configuration is as described in the preceding chapter. The selection and design of specific components is described below, using the analysis from the previous chapter.

7.2.1 Source actuator

The design requires a low-impedance source actuator. A Copley Thrusttube 2504 [29] tubular linear motor was used. This actuator can provide continuous force up to 51 N (though the present current amplifier may limit the continuous force to around 31 N) and peak force of 281 N (80 N with present amplifier). It has position-dependent static friction between 3 and 4 N in magnitude and endpoint inertia, due to the mass of the forcer, of 1.45 kg [86]. With proportional motion control, it can produce stiffness of at least 50 kN/m. This is the same model motor used for the direct-drive robot module described in example 2.2, though the bearings and packaging differ slightly.

7.2.2 Piston-cylinder and moving seals

Selecting an appropriate hydraulic cylinder, with appropriate seal type, is critical to meeting the low-mass and low-impedance system requirements. The requirements for the cylinder in this design are that it have low mass (the total endpoint mass budget is 2 kg), effectively contain the closed fluid, convert from pressure to force (and vice versa), and that it do this with an absolute minimum of friction. Stiction is of particular concern, and should be kept to less than 2 N. For some of the most common types of seals, the low stiction and low leakage requirements are in strong conflict.

The cylinders must be able to withstand the pressures needed to transmit the
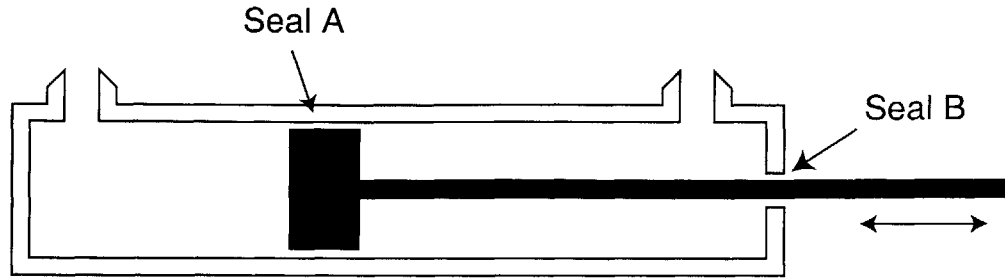
Seal A

Seal B

Figure 7-1: Two possible areas requiring seals in a hydraulic cylinder. Seal "A" is across the piston, and seal "B" is at the piston rod exit.

desired force. If the allowable pressure is large, a cylinder smaller in diameter can transmit the same load while taking up less space and probably having less mass. The application calls for forces up to at least 65 N. Any hydraulic cylinder, and in fact all parts of the hydraulic system, must be able to withstand at least $P_{max} = \frac{65N}{A_{pist}}$, where A_{pist} is the piston area and P_{max} is the maximum pressure. If a neutral bias pressure is used, the operating pressure may be twice this, or more.

Finally, the cylinders and seals must be compatible with the selected working fluid. Working fluids considered included petroleum-based hydraulic oils, water, and water-based hydraulic fluid, and are discussed below.

The seal is the most critical component of the cylinder as it dictates both fluid containment and friction. The piston moves within the cylinder, and this interface, labeled "A" in figure 7-1, must be sealed. If dual fluid paths are used, the interface between piston rod and cylinder exit, labeled "B" in figure 7-1, must also be sealed, and is perhaps the more important of the two. If fluid crosses the piston seal, the relationship between the position of the two pistons may change; if fluid crosses the exit seal and escapes the constant volume circuit, it is replaced by air, leading to compliance in the lines. The first problem introduces inaccuracy in motion regulation; the second can introduce serious dynamic flaws.
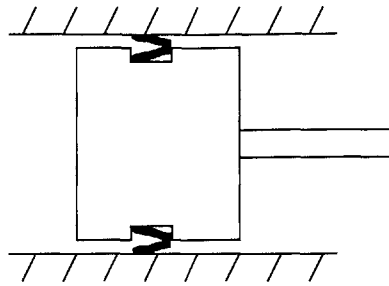
Figure 7-2: Cross section of a piston with a lipseal.

Sliding contact seals

The seals most commonly used for industrial pneumatic and hydraulic applications are sliding contact seals. Typically made of rubber or another elastomeric material, the seal rides on one surface and brushes against the other, compressed slightly to push against both surfaces. Many varieties of such seal designs exist, ranging from simple O-rings to a number of more complex geometries. As a rule, the greater the force with which the seal pushes against the surfaces, the better the seal. Coulomb's law dictates, however, that this larger normal force results in more sliding friction. The coefficient of friction can sometimes be reduced by proper material selection and lubrication, but in general tighter sliding seals produce more friction.

One common type of sliding seal is known as a lip seal and is sketched in figure 7-2. Because there is a space in the "cup" of the seal, pressure increases inside the cup result in increased forces pushing the seal against the surfaces. The result is a seal that gets tighter as the pressure gets higher, which is when the seal needs to be tighter to control leakage. Somewhat lower contact forces, and therefore lower friction, can be preserved at lower pressures.

An analysis was performed to determine whether the cylinder diameter influences the ratio of output force to friction force with seals of this type. This could guide a decision as to whether, aside from pressure and force considerations, either smaller or larger cylinders provide advantages. If a lip seal is cut, unwrapped and laid flat, it looks like figure 7-3. The length is equal to π times the piston diameter, and l is the width of the lip portion of the seal. These two parameters define the projected area

271
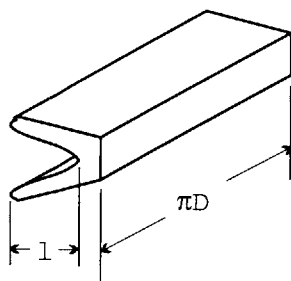
Figure 7-3: A lipseal, cut and unwrapped.

on which pressure acts, assumed to be the same area that comes into contact with the cylinder and produces friction.

Assuming a simple Coulomb friction model, the friction force is defined as $F_f = \mu N$ where μ is the sliding friction coefficient and N is the normal force. The normal force results from pressure on the interior of the seal (between the two lips) applied to the projected area of the lip, leading to the result:

$$F_f = \mu P l \pi D \tag{7.1}$$

The nominal output force is equal to the pressure P times the area of the piston:

$$F_{out} = P\pi \frac{D^2}{4} \tag{7.2}$$

Dividing the previous two equations gives the ratio of friction force to the output force:

$$\frac{F_f}{F_{out}} = 4\mu \frac{l}{D} \tag{7.3}$$

To improve performance, this quantity should be minimized. Since μ is determined only by the materials in contact, this means the quantity $\frac{l}{D}$ should be minimized. This quantity is the aspect ratio of the seal, however, and has no fundamental reason to change as D is varied. If l is reduced, the leakage path length is also reduced, and the leakage rate increases as friction decreases. This demonstrates one fundamental seal tradeoff between friction and leakage.

272
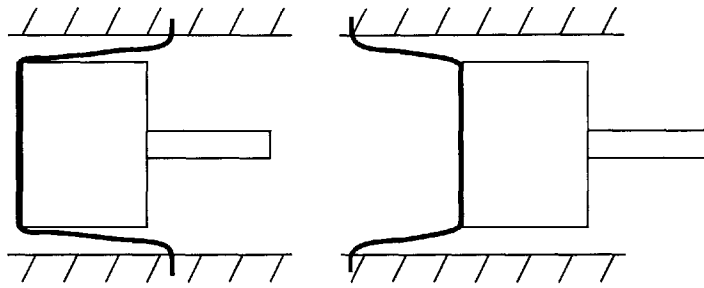
Figure 7-4: A rolling diaphragm-sealed piston, in its two extreme positions.

This analysis indicates that there is no obvious benefit gained, in terms of reducing relative friction, by using either smaller or larger pistons with sliding lip seals. This is because the output force, which scales with the piston area $\pi \frac{D^2}{4}$, and the friction force, which scales with the circumferential area lD, both scale quadratically with cylinder diameter, assuming that l scales linearly with D. Rather than changing diameter, the best way to reduce seal leakage is by increasing normal force and therefore friction.

Rolling diaphragm seals

A dramatically different type of seal is the rolling diaphragm seal. Shown schematically in figure 7-4, the rolling diaphragm seal is clamped into the outside wall of the cylinder around its outer diameter, and is affixed to the piston head. The diaphragm is made of a highly flexible elastomer. There is a significant clearance between the piston and cylinder, and in this space, the diaphragm rolls up or extends, depending on the relative position of the two parts.

The rolling diaphragm design creates an absolute seal between the two fluids; they in fact reside in two separate chambers. Manufacturers also claim that they have zero breakaway friction and zero spring rate (see for example [33]).

The main drawback of rolling diaphragm seals is that the ratio of stroke to diameter is sharply limited. For off-the-shelf seals, the best ratio of stroke to diameter appears to be about 1.5. Thus if a 9-inch stroke is desired, a cylinder at least 6 inches in diameter is needed.

Because it promises to completely eliminate leakage, the rolling diaphragm seal

273
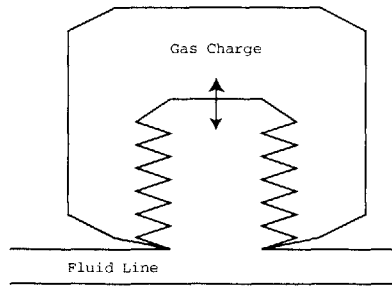
Figure 7-5: A bellows accumulator.

appears to be the best option for short stroke applications, or for any application where exceptionally large diameter cylinders can be tolerated. For applications with large stroke length, however, and where bulky cylinders are not acceptable, this seal type is not effective.

Bellows designs

Bellows, often made of metal, are used in hydraulic systems as accumulators. The bellows offer an expandable volume of contained fluid that typically expands or contracts against a gas charge in a closed volume. This arrangement is shown in Figure 7-5.

Bellows could similarly be used in a hydraulic cylinder, with the inside of the bellows representing one fluid chamber and the outside the other. Such a configuration should have very low friction and stiction. Depending on the specific implementation, bellows are expected to have some substantial stiffness about a preferred position. It is also possible that the bellows could buckle under high pressures if not properly constrained. Achievable ratios of travel to total length are not known, but are likely to be small. No commercially available bellows-type cylinders could be located.

"Seal-less" piston-cylinders

Another anti-stiction and low-friction family of cylinders consists of those that use no contacting seal at all. In the embodiment manufactured by Airpot Corporation, a precision ground graphite piston slides into an ultra high-precision pyrex glass tube
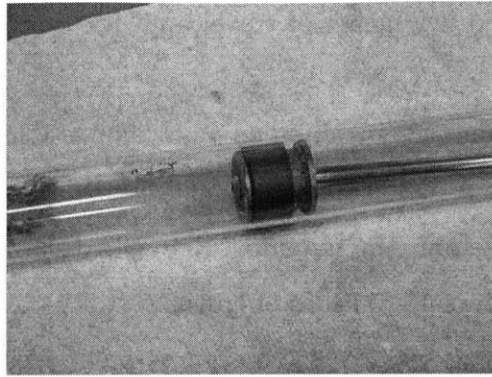
Figure 7-6: The pyrex glass cylinder and sintered graphite piston of an Airpot.

[2]. The tube is manufactured by heating a precision glass tube, placing it over an ultra high-precision rod, drawing vacuum between the two and sucking the glass onto the rod exactly. The inner diameter of the glass tube and the outer diameter of the piston are measured with an air gauge, and matched to within a few ten-thousandths of an inch. Typical clerances range from 0.0003" to 0.001". A photograph of a piston/cylinder set is shown in figure 7-6.

This design, known by its trade-name Airpot, is intended to be a pneumatic cylinder or dashpot. Published leakage rates are somewhat larger at lower pressures than they are for comparably sized sliding seal cylinders, but the friction reduction is dramatic. The manufacturer specifies that pressure differentials as low as 0.2 psi will move the piston in the cylinder. The published pressure limit is 100 psi, lower than many sliding seal cylinders which in some cases go into the thousands of pounds per square inch. For the output force levels expected in this design, breakaway force of less than 0.5 N is expected.

The manufacturer expressed concern that anti-stiction behavior would be lost if fluid other than air was used in the cylinder; however, no problem is evident from the basic physics. Neither pyrex nor graphite are expected to corrode significantly or absorb much of the proposed working fluids.

A similar seal-less design can be found in glass syringes, once common in medicine. These use matched parts that are lapped to fit each other. Commercially available

glass syringes have only a few inches of travel, and are limited to approximately 30 psi pressure.

In the interests of minimizing friction as much as possible, Airpot cylinders were chosen for the prototype. A piston diameter of 0.627 inch was selected. The maximum cylinder length available from the manufacturer with this diameter was 14 inches. This permits a maximum of 13.5 inches of travel.

With the noncontacting seal design, friction is minimized. This is only possible because the clearance between piston and cylinder is extremely small. Sealing at the piston rod, seal B in figure 7-1, presents another challenge if dual fluid paths are used to provide bilateral actuation. Seal B is not available in the same materials or the same level of precision as seal A. If a sliding seal is used (the most obvious solution), this adds friction and the benefits of seal A's low friction are obscured. Rolling diaphragm and bellows seals present as much of a problem at seal B as at seal A, as the total range of travel at the two seals is identical. Because of the difficulties in sealing at the cylinder exit with low friction (verified in testing with Airpot pre-fabricated bilateral cylinders), the bias pressure technique was selected for this prototype. Details of this implementation are provided below.

7.2.3 Hydraulic fluids

Fluid requirements vary across hydraulic applications. Factors to consider include compressibility, viscosity, density, lubricity, corrosion, cost, environmental factors, and flammability.

For this application, it is desired that the fluid have a low compressibility and low viscosity in order to meet the stiffness and damping requirements for the system, and be compatible with the hardware selected (not corrosive). Additionally, a fluid with lower density is preferable, in the interests of low mass and endpoint inertia. Lubricity, which is of prime concern in hydraulic machinery where the working fluid also must lubricate mechanical bearings, is less important here, as the architecture does not include any parts in rolling or sliding contact (except the seals themselves, which do not require lubrication).

276

	Hydraulic Oils	Water	Water-Glycol
Weight	0.9 * water	-	1.1 * water
Viscosity	30-100 * water	-	comparable to oil
Lubricity	good	poor	moderate
Corrosion	good	poor(use proper materials)	moderate

Table 7.1: Comparison of major fluid properties.

A range of hydraulic oils, water, and water-based hydraulic fluids were considered. Air and other gaseous fluids were ruled out because of their high compressibility. Table 7.1 shows a comparison of ISO hydraulic oils, water, and water-glycol, one representative water-based hydraulic fluid. Generally, oil is slightly lighter than water and significantly more viscous. Water-glycol has many of the lubricity and corrosion properties of oil and is slightly heavier than water. Water is by far the cheapest, least viscous, and most environment-friendly. It corrodes metals but proper material selection eliminates this problem. It is a poor lubricant, but this is not expected to be a factor in this design. Water presents the possibility of bacteria growth over time, but when this is an issue it can be "poisoned" with some amount of alcohol. Thanks mainly to its low viscosity, abundance, and cleanliness, distilled water was used.

7.2.4 Fittings and static seals

The connection between the rigid cylinder body and the flexible hose turned out to be one of the most problematic sources of undesired fluid damping and energy loss. Standard pneumatic and hydraulic screw fittings use a threaded metal piece that is crimped on to the hose and mates with a thread on the cylinder. The piece that is crimped to the hose significantly reduces the effective diameter, causing large head losses. An alternative connection is needed that does not drain so much energy.

To reduce line impedance, fittings are required that can connect the flexible hose to the rigid structure with minimal obstruction of the fluid path. An effective fitting must also provide a tight static seal so that the fluid does not leak under pressure. These objectives can be achieved at moderate pressures (to at least 150 psi) with quick-disconnect fittings that attach and seal at the outer diameter of the hose. A
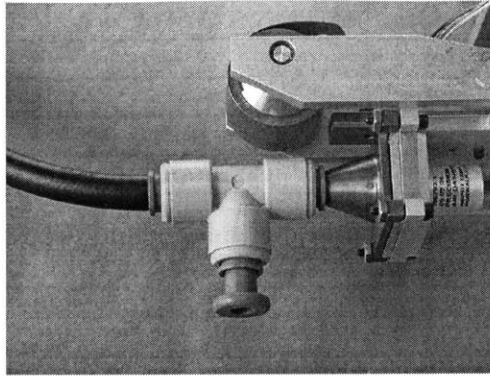
Figure 7-7: Endbell and quick-disconnect fittings.

hose of a specific outer diameter is inserted into the fitting and squeezes past a collar with wire teeth and an O-ring. After the hose is inserted, if it is pulled out the collar moves with it and the internal geometry of the fitting squeezes the wire teeth into the hose, preventing it from being removed. The collar must be separately held in place to disassemble the connection. The O-ring remains tightly around the outer hose diameter, preserving a tight seal. These fittings have almost no internal obstruction; for example, the size used with hoses with 1/2 inch outer diameter has a minimum internal diameter of 3/8 inch. Therefore any contribution of extra head loss is minimal. The fittings only work if the wire teeth can indent the hose material, and if the hose material is a suitable diameter. This accomplishes the task of attaching the rigid fitting to a flexible hose; the task remains to connect the fitting to the hydraulic cylinder. This is accomplished with an endbell made of soft copper, so that the fitting could grip its outer diameter. The design of the interior of the endbell is discussed below in the section on fluid mechanics. The endbell, fitting, and hose assembly is shown in figure 7-7. The primary drawback of this fitting design is that the largest available size uses hoses with 1/2 inch outer diameter. Thus only relatively small hoses can be used. Small hoses are good for flexibility and integration into a robot, but poor for reducing inertia and damping. The implications of this are discussed in the section on fluid mechanics.

The fitting provides a static fluid seal between the endbell and the fluid line.

278
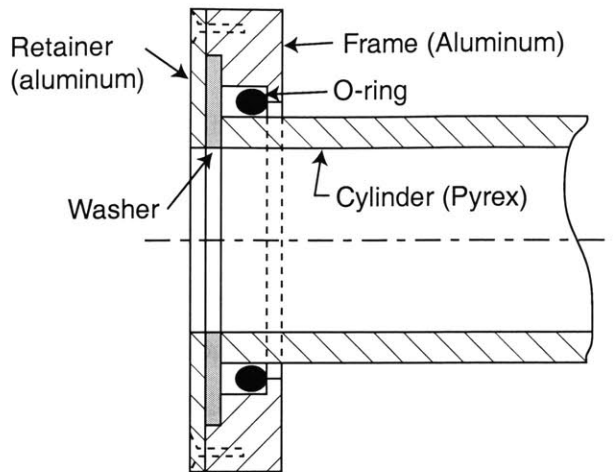
Figure 7-8: Cross-sectional drawing of static seal at cylinder end. An O-ring seals the cylinder end and provides radial cushioning. A washer provides axial cushioning.
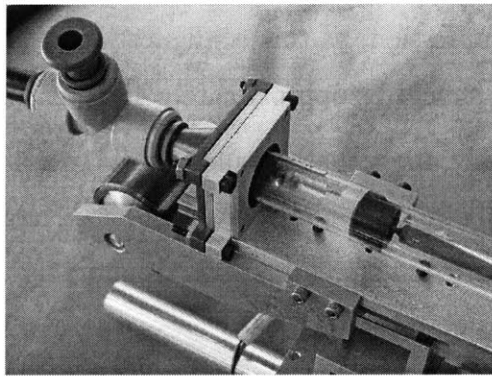
Figure 7-9: Endbell, frame, and cylinder statically-sealed assembly.

A quality seal must also be established between the cylinder and the endbell. The cylinder is simply a glass tube. Each end of it is mounted with an O-ring around its radial exterior to an end frame piece. The O-ring provides a seal as well as a mechanical cushion to protect the glass from impact and vibration. Each end of the glass tube is mounted axially with a rubber washer between it and the frame. The washer is slightly compressed to hold the cylinder in place. The washer, O-ring, cylinder, and frame end are shown in figure 7-8. The frame is bolted to the copper endbell. Another washer is compressed between the two pieces to seal in the fluid. The assembly of cylinder, frame, and endbell is shown in figure 7-9.

279

7.2.5 Bias pressure mechanism

To avoid the need for extra sliding seals away from the piston itself (seal B in figure 7-1), the prototype uses a single fluid chamber to apply forces in both directions. The neutral bias pressure method described in section 6.3.4 is used. A reliable mechanism is needed to consistently apply bias forces at each piston. Because the piston areas are equal, the forces are equal. For optimal performance with simple control algorithms, the forces should be as close to constant as possible under all operating conditions. Thus the system that applies the forces should have minimal dynamics.

To keep the fluid pressure above atmospheric pressure for all system forces, the bias force should be greater than or equal to the maximum tension force. This application requires up to 45 N downward force, and up to 65 N upward force. The higher force should be applied in compression, as this configuration requires a lesser bias force. Thus the bias force F_{bias} should be approximately 45 N, or 10 lb.

As mentioned previously, the actuator could be used to apply bias force on one side of the transmission, simply by adding F_{bias} to every force command. The Thrusttube 2504, however, has a continuous force limit of around 51 N. Clearly an additional 45 N applied at all times would dramatically reduce actuator authority, requiring a larger source actuator.

On the outer side of the transmission, some bias force could be applied using the weight of the robot handle and attached hardware, or by adding additional weight. This approach has two disadvantages, both related to the inertia. The inertia would prevent the applied force from being exactly correct whenever the endpoint acceleration is nonzero, degrading performance. Furthermore, the magnitude of the mass that would be required for this, 4.5 kg, clearly prohibits this approach as this would ruin the specifications both on endpoint inertia and on the actuator mass.

An alternative method of applying nearly constant forces is with constant-force springs. These springs are made of flat spring steel that is rolled into a coil. As the coil is unwrapped, a certain portion of the steel develops stresses that tend to pull the unwrapped portion back around the coil. The region that generates the
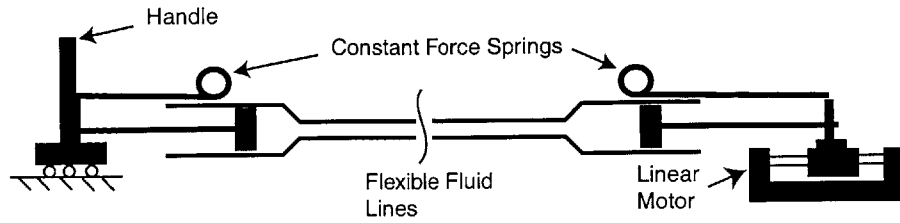
Figure 7-10: Schematic of a two-piston system with constant-force springs applying bias pressure.

relevant stresses is, in theory, identicallly sized regardless of deflection, and therefore the force applied is theoretically constant with position. The springs have minimal inertia and therefore are nearly static. In practice, constant-force springs do not apply perfectly constant force, and they can be difficult to work with, mostly because their energy storage makes them difficult to attach and detach. However, they offer one purely mechanical means of applying a theoretically constant bias force without adding substantial mass or dynamics.

Constant-force springs were selected to apply F_{bias} in the second prototype. Springs that apply a nominal force of 47.1 N were chosen. These springs each have an outer diameter of 3.9 cm and a mass of 82 g. The springs were configured to uncoil as each piston extends, while applying force on the piston back toward its cylinder. A schematic depiction of the fluid transmission system with the constant force springs is shown in figure 7-10.

Each spring is mounted to a drum that is free to rotate within low-friction ball bearings. Figure 7-11 shows a photo of the constant-force spring mounted to the outer cylinder.

7.2.6 Fluid mechanics and hose selection

The fluid lines convey both fluid volume and force between the two cylinders. Their geometry is critical to the endpoint impedance, and is modeled with the fluid mechanics model given in Chapter 6. The model specifies a minimum inner line diameter to meet the inertia and damping specifications, but these parameters must be bal-

281
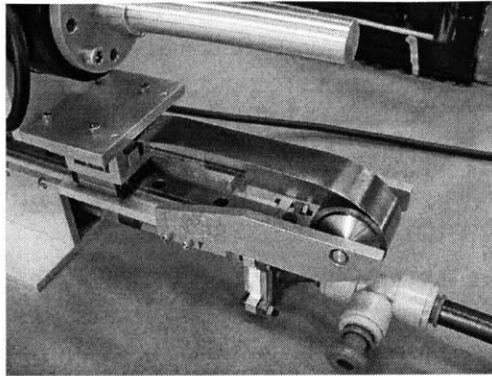
Figure 7-11: Constant-force spring mounted to endpoint module of fluid actuation system, to provide bias pressure.

anced against several practical concerns. Both inertia and damping are dramatically reduced as the line diameter is increased, but this can run contrary to the additional requirements that the hose be flexible in bending, radially stiff, and capable of bearing sufficient fluid pressure. Diameter should also be minimized in order to keep the lines lightweight and low-profile. One aggressive target is for the lines to be comparable in size and flexibility to shielded electrical power cables used to carry 5-10 amp to electromechanical actuators. As noted above, to use the selected fitting hoses are restricted to 1/2 inch or smaller outer diameter.

Fluid line types

Several different hose types were considered for the first prototype, including flexible metal, rubber, and thermoplastic, the latter two with and without fiber reinforcement. Metal braided hoses offer the greatest resistance to hose damage, but are less flexible than the other designs. Rubber hoses tend to be the most flexible, but require reinforcement to satisfy the pressure requirements. One promising class of hoses conforms to the SAE 100R7 specification. This type of hose with a nylon core, fiber reinforcement and a polyurethane cover is readily available in inner diameters as small as 1/8 inch and can withstand at least 2500 psi. The minimum bend radius on the 1/8 inch hose is 1/2 inch, by far the smallest of the hoses considered. This hose has an outer diameter of 0.34 inch. Larger versions of the same hose are also available.

282

Because of the high pressure rating, the walls of the SAE100R7 hoses are thick, so the inner diameters are small and result in high viscous damping and inertia. Another promising reinforced hose conforms to SAE30R3, is made of chlorinated polyethylene with a nitrile cover, and is intended for fuel transfer. One size of this hose has a 5/16 inch inner diameter, a 1/2 inch outer diameter, and can bear 300 psi. This hose is extremely flexible. An alternative is nylon tubing, with an outer diameter of 1/2 inch, an inner diameter of 3/8 inch, and a pressure rating of 285 psi. This provides the maximum inner diameter that can work with the 1/2 inch fittings and still withstand the operating pressures, but is more rigid than the rubber hoses. The minimum bend radius of this hose is 5 inches.

The fluid mechanics calculations below are done with these practical limitations and possibilities in mind. The implications of the calculations for the practical design are noted where relevant.

Damping

Practical experience with therapy robots suggests that a linear damping coefficient of less than 10 Ns/m is acceptable and does not interfere substantially with low-force attempts to move. Analysis of the fluid mechanics shows that the flow in the fluid lines is likely to be turbulent at speeds within the specified range for the robot, and therefore that the damping is likely to be nonlinear. For instance, if the hose is a straight pipe with a 3/8 inch inner diameter, and one of the pistons (as sized in section 7.2.2) is moved at a speed of 0.15 m/s, the Reynolds number in the line is 4000, while turbulence is likely to occur above $Re = 2300$. When the hose is not perfectly straight and the turbulating effects of the fittings are included, the onset of turbulence is expected at even lower speeds.

The module was designed to operate up to speeds of at least 1 m/s. The most conservative interpretation of the specification is to guarantee that the force is below the force that would be produced by a 10 Ns/m damper at all speeds in the operating range, meaning that a friction force of 10 N at 1 m/s should be the maximum. Because turbulent internal flow can produce a highly nonlinear damping curve, however, this

may be extremely difficult to achieve. As the specification of 1 m/s was selected mainly from estimation of maximum patient arm speeds, it is worth examining how fast similar robots typically are used in practice, in order to interpret the specification in the fairest way possible.

Using a set of data from stroke patients engaged in planar arm therapy, patient arm movement speeds across a number of trials were examined. The data came from 64 different outpatients performing a total of over 52,000 movements throughout their course of therapy. Because the patients are outpatients, they are at a higher level and therefore expected to be capable of faster movements than inpatients. For each individual movement, the peak speed, or the maximum speed reached at any point in the movement, was computed. A histogram of these peak speeds, placed into 100 equal bins, is plotted in figure 7-12. The mean peak speed across all movements is 0.156 m/s. The standard deviation of the peak speeds is 0.074 m/s. The maximum speed achieved in any of the trials considered here is 1.065 m/s. In over 99% of all trials, the speed stayed below 0.4 m/s. Thus it is critical that the damping specification be met below speeds of approximately 0.4 m/s, but somewhat higher damping is permissible at higher speeds, because it is highly unlikely that the robot will be operated at these speeds regularly. Note that there is no safety risk in higher damping at higher speeds; only the desired impedance is compromised.

Given the practical restriction on the outer hose diameter, only a small range of inner hose diameters are considered. Using the fluid mechanics model from the previous chapter, the expected major damping due to the fluid lines is plotted in figure 7-13 for inner diameters of 1/4 inch, 5/16 inch, and 3/8 inch. A 3 m fluid length is assumed, and the 10 Ns/m target is also plotted. With these parameters, transition from laminar to turbulent flow is expected at a piston velocity of 0.057 m/sec for the 1/4 inch hose, 0.071 m/sec for the 5/16 inch hose, and 0.086 m/sec for the 3/8 inch hose. Practically, transition likely occurs at somewhat lower speeds because the hose is not a perfectly straight pipe. From figure 7-13, a hose with inner diameter greater than or equal to 5/16 inch is expected to meet the damping specifications.
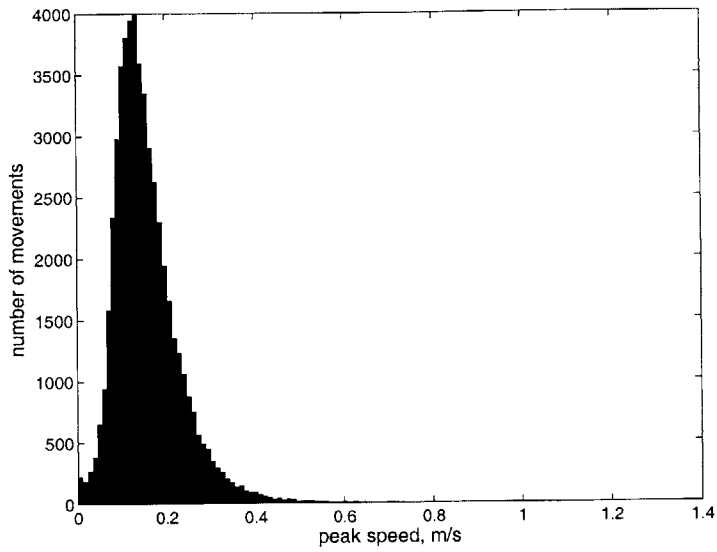
Figure 7-12: Histogram of stroke patient peak speeds across 52,000 movements.
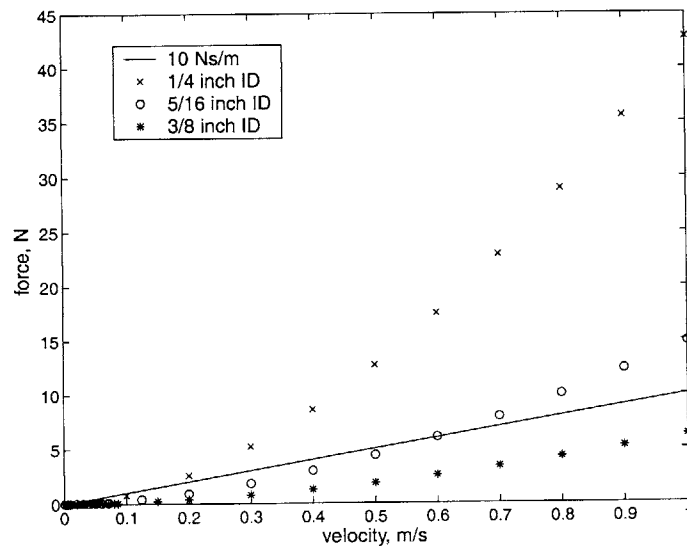
Figure 7-13: Predicted major damping due to fluid lines for several different line diameters.
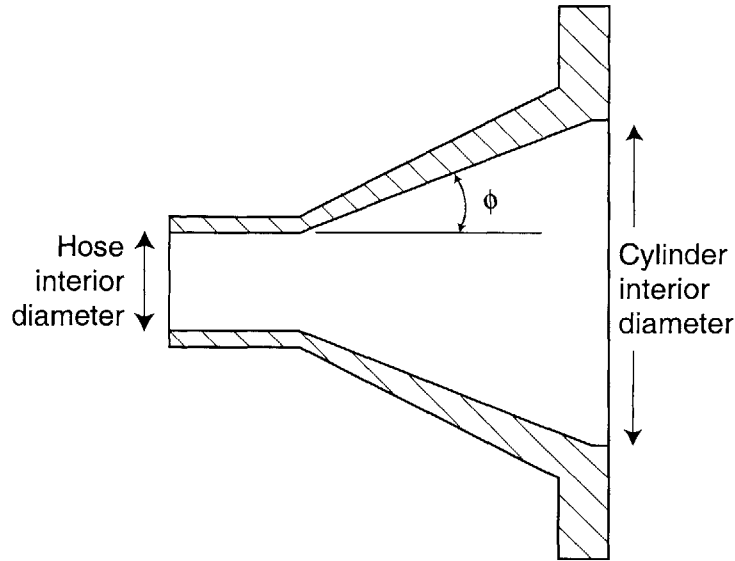
Figure 7-14: Cross-section of interior cylinder endbell geometry. The gradual change in diameter reduces losses.

The interior hose diameter is smaller than the interior piston diameter. The cylinder design must accommodate the transitions from large to small and small to large diameter, both of which can result in substantial losses and contribute significantly more friction. Expansion to a larger flow diameter contributes an equal or greater amount of loss than contraction to a smaller diameter. In figure 6-8, the transitions are depicted as sharp corners, but a more gradual transition can reduce losses significantly. Designed with the goal of adding less than 5% additional friction at all speeds in the operating range, a linear change in diameter, characterized by the angle ϕ, was selected and is shown in figure 7-14.

Because the fittings have an internal diameter of 3/8 inch, when 5/16 inch hose is used, an abrupt transition between these two diameters is required at both ends, where the hose meets the fitting. Thus the cylinder endbell is designed to transition from the 0.627 inch cylinder diameter to 3/8 inch (an area ratio of $AR = 2.8$). From figure 8.17 in [44], with an area ratio of approximately 0.7, the contraction and expansion between the hose and fitting each contribute $K = 0.1$ (where minor head loss is determined by equation 6.10). This increases the fluid system friction force by approximately 1.8% at 0.5 m/sec, and 2.2% at 1 m/sec. Thus the nozzle and diffuser

due to the endbell should be designed to add less than approximately 3% to the total friction. Working backward from the major loss calculations, this should be less than 0.135 N at 0.5 m/sec (piston velocity), and less than 0.45 N at 1.0 m/sec. This second point is more restrictive, and dictates that the total loss coefficient $K < 0.27$ to meet this specification. When ϕ is between 25° and 30°, K for the nozzle is less than 0.07, so most of the losses come from the diffuser (see table 8.2 in [44]). Computing the loss in a diffuser is more complicated, and depends on a number of variables. For fully developed turbulent flow, some data is provided in terms of the *pressure loss coefficient* C_p in figure 8.18 in [44]. C_p can be related to K through the relationship:

$$K = C_{pi} - C_p \tag{7.4}$$

where C_{pi} is the ideal pressure loss coefficient, determined by:

$$C_{pi} = 1 - \frac{1}{AR^2} \tag{7.5}$$

For this example, $C_{pi} = 0.9394$. So, to get $K \approx 0.25$, a C_p around 0.7 is needed. From the data in [44], this requires ϕ between 2.5° and 5°. Because this would produce an extremely long endbell, an angle of $\phi = 10°$ was selected. This results in a pressure recovery coefficient of 0.45 and $K \approx 0.5$. At a piston speed of 1 m/sec, the minor losses add roughly 8% to the major friction losses. This percentage steadily decreases at low speeds. With $\phi = 10°$, the length of the endbell expansion is 0.715 inch.

Inertia

The target inertia for the system, 2 kg, is difficult to meet because the source actuator itself has inertia of 1.45 kg. This limits the fluid system to contribute only 0.55 kg of additional inertia. To meet this specification with a line length of 3 m, the inner line diameter must be at least 0.697 inch, larger than the piston-cylinder diameter. Because the outer hose diameter is limited to 0.5 inch, this is clearly not possible. Inertia, not damping, is the limiting factor in this design. The maximum practical

inner line diameter should be selected. The nylon hose, with an inner diameter of 3/8 inch, should result in fluid inertia of approximately 1.77 kg. The SAE30R3 fuel line, with an inner diameter of 5/16 inch, should produce an apparent fluid inertia of approximately 2.52 kg. (Each of these includes 0.07 kg from the fluid in the pistons). Because the fuel line is considerably more flexible, it was selected. To meet the inertia specification while retaining the piston-cylinder diameter and the desired line length and level of flexibility, a larger hose diameter is required. This could be used if larger fittings are custom-made that use the same design, or if alternative fittings that similarly avoid constricting the fluid path are used. Of course, using a significantly larger line diameter would dramatically reduce damping, and thus a more constrictive fitting design may be acceptable without violating the damping specification.

Final hose design

For this prototype, an SAE30R3 fuel hose was selected, with an outer diameter of 0.5 inch and an inner diameter of 5/16 inch. This hose is extremely pliant and flexible compared to other available hoses of similar dimension or to electrical power cables of similar size. The expected friction due to this hose, when completely straight and level, is shown in figure 7-13, and the expected fluid inertia is 2.52 kg. The hose is radially reinforced, so an estimate of line compliance is difficult; this was determined through testing.

7.2.7 Linear bearing and structure

The vertical module must bear a downward vertical load of up to 65 N and side loads of up to 45 N in all directions, applied at the handle. In order to reduce the overall length of the module (with piston retracted), and to therefore reduce the structure needed, the piston rod was made of three pieces so that it could be "folded" over the top edge of the cylinder, and the handle was placed parallel to the cylinder. This is depicted in the solid model in Figure 7-15. Because the handle is off-axis from the piston, any forces applied by or to the module produce torques.

288
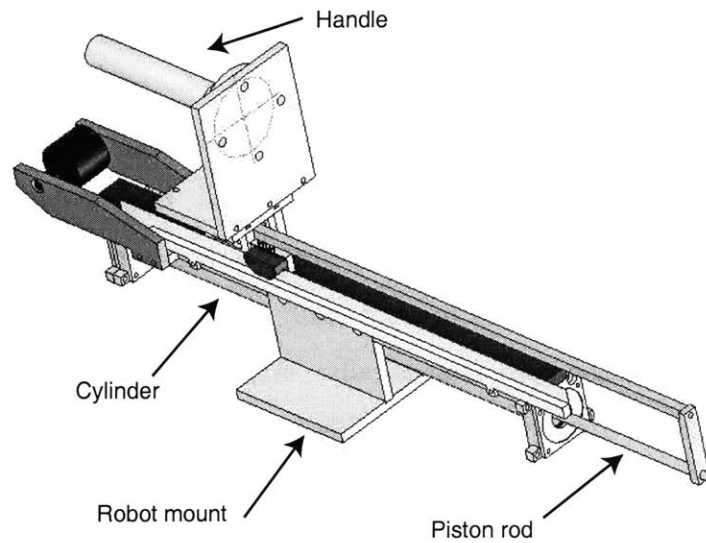
Figure 7-15: Solid model of endpoint assembly.

All of these forces and torques must be isolated from the piston-cylinder interface, so as not to affect the friction properties of the seal. An ideal linear bearing for this application would bear torque in all three axes as needed, and force in two axes, while permitting motion without friction in the primary direction of travel. Given the mass sensitivity of the application, it is also critical that the bearing size and mass be minimal.

A low-profile ball slide was chosen to provide 5 axes of support and 1 axis of travel. For bearings of this style, the torque requirements of this application dominate the force requirements; typical force capacities of bearings that sufficiently support the torque requirements are in the kilo-Newtons. (Using multiple smaller bearings spaced appropriately is likely to produce a more mass-optimal design). The worst-case torsional loading condition occurs when the module provides the full 65 N of upward load. In this case, both the human subject and the constant-force spring pull down, so fluid pressure must provide 112 N of force at the piston (65 N to offset the load plus 47 N to offset the constant-force spring). Assuming the piston and spring forces are each located .025 m from the bearing axis, and the environment force is located 0.1 m from the bearing axis, the total torque on the bearing is 10.5 Nm. By

similar reasoning, the peak torque on the bearing in the other axes is expected to be less than 4.5 Nm.

Initially, a ball slide with a single row of balls on each side of the rail was selected. In order to support forces and torques in all 5 axes, such a design requires each row of balls to contact nominally at two points per ball using a Gothic arch groove design (see [89]). Although this bearing easily supported all applied loads, under torque loading its friction rose dramatically. With the fluid under a bias pressure, there is a substantial torque on the bearing even when the endpoint force is near zero. But this load condition is precisely when it is most critical that friction be minimized. Because the friction of the first bearing was unacceptable, it was replaced with a bearing that featured two rows of balls per side. In such a design, each row of balls must only contact the rail, nominally, in one spot. This change dramatically improved low-friction performance. The final design uses a THK SHW12CR ball slide that provides a minimum of 22.8 Nm of support in torsion, providing a design factor of more than two.

The three-part piston rod must also bear the large forces transmitted between the piston and the handle/bearing assembly. The rod piece that attaches to the piston directly is primarily limited by buckling, as it has ball joints on both ends to protect it from bending moments. It is a cylindrical steel rod with 1/4 inch diameter. The rod piece that mounts to the bearing carriage must support loads in tension/compression and bending. The deflection due to bending turned out to be the limiting factor in its design. The 112 N maximum load at a distance of 4.3 cm from the rod axis produces a maximum bending moment of 4.6 Nm. The effective rod length from its secure mounting on the bearing carriage to the cross piece is approximately 0.38 m. Applying the moment-area method (for beams in bending) to determine deflection [55], if the aluminum rod has a rectangular cross section of 3/8 inch (0.95 cm) by 1/2 inch (1.27 cm), deflection under the maximum load should be less than 2.8 mm. This is perhaps larger than desired, but small enough to keep the rod clear of the cylinder rim, and the ball joints help the design to tolerate such deflection. The cross piece of the piston rod is subject to the same torque, but because its effective length is much

shorter, its deflection is minimal.

Similarly, the handle assembly that mounts to the bearing carriage was sized for the endpoint forces and torques. The rest of the structure bears minimal loading. The steel bearing rail prevents bending of the structure under side loads. The pressure requirement is relatively low, so the parts around the endbell that contain fluid are substantially oversized for this loading.

7.2.8 Leakage and fluid makeup

Because they use non-contacting seals, Airpot piston-cylinder sets have a finite air gap between the piston and the cylinder. The manufacturer estimates the size of this gap to be between 0.0003 and 0.001 inch (diametrically). Modeling leakage flow as in section 6.3.5, an upper and lower bound on leakage can be estimated. Using the bias pressure of the system at rest to determine the average leakage, the leakage flow is determined by equation 6.29:

$$Q = \frac{\pi D a^3 \Delta P}{12 \mu L} \tag{7.6}$$

The gap a is measured radially, and is taken to be half of the diametric estimates given above. Using the minimum diameter and doubling the total to account for the two pistons, the leakage flow estimate is 3.04×10^{-5} cm^3/hr (for both pistons), which produces 15.2 cm/hr of piston travel. Using the maximum diameter, the estimate is 1.13×10^{-3} m^3/hr, which produces 566 cm/hr of piston travel.

Because the expected leakage flow is large, a mechanism to replace the fluid as it is lost is essential. Ideally, this makeup system should supply fluid equal in volume to the lost fluid. The fluid should be supplied at high impedance from a flow source to avoid any impact on the transmission pressure, which is dictated by the forces applied at the two pistons. Because the leakage rate depends on the real pressure in the fluid adjacent to each piston, it varies as forces are applied and as the system moves. An ideal makeup system would adjust the input flow in real time to keep the fluid volume in the transmission constant. The leakage could be estimated in real

time using a variety of inputs, including position measurement from both pistons, force measurement at the end-effector, force command at the actuator, and possibly a model of the transmission.

The fluid makeup system requires an independent flow source. Ideally this should be a positive-displacement pump that can supply a specified and controllable fluid volume. A pump with smooth delivery, like a syringe pump, is preferable to a cycling pump because variations in flow input might introduce undesired vibration in the fluid transmission.

In order to implement a low-cost makeup system quickly, a simple feedforward pump system was implemented to approximate the performance of a more complex makeup system. A Pulsafeeder diaphragm pump, shown in figure 7-16, was used as a source. This pump is manually adjustable with nominal flow rate from 0.1 to 10 gallons per day (leakage calculations translate to 0.15 to 3.6 gallons per day). It begins a release phase approximately once every 2.5 seconds. The relatively low supply frequency and low supply rate, resulting in very small energy input from the pump, make the makeup flow nearly imperceptible to humans at the haptic interface during transmission operation. In the absence of a leakage estimator or a mechanism for precisely metering the makeup flow, the system was manually set to provide sufficient flow to counter the average leak rate. This was done by observing the system at rest with no applied forces at either piston, and adjusting the input flow so that the net relative movement over time between the two pistons was negligible. This makeup system functions adequately for system testing and operation, though it is susceptible to drift if forces are applied primarily in one direction over a long period of time (because the resupply rate is fixed).

7.3   Testing and characterization

The outer cylinder assembly of the prototype is shown in figure 7-17. The source actuator and motor-side cylinder assembly are shown in figure 7-18. The system was controlled with a PC running RT-Linux (Realtime Linux). Force was measured
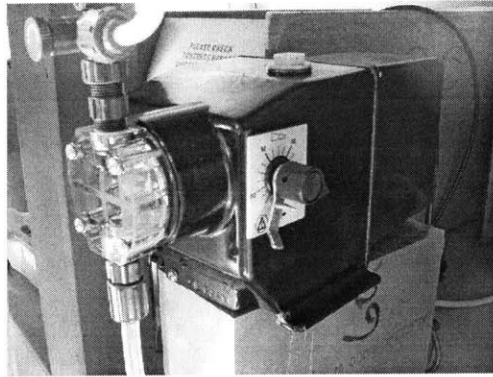
Figure 7-16: Diaphragm pump for fluid makeup.

at the interface between the transmission and the environment, with a six-axis ATI Gamma transducer mounted at the handle. Motion was measured both at the linear motor encoder and at the encoder mounted on the cylinder, but only the linear motor encoder was used for feedback control. When zero impedance was desired, the actuator sent no command. When nonzero impedance was desired, simple impedance control was used to implement it. A series of tests were run on the system to determine its force and impedance capabilities.

7.3.1 Weight of outer assembly

The mass of the outer cylinder assembly (here referred to as the "weight" to avoid confusion with the apparent endpoint inertia) is critical as it strongly influences the endpoint inertia, the loads on the robot structure and, in some robot configurations, the gravity loads borne by both the actuators and the subject.

The total weight of the outer cylinder package, not including the hose or the working fluid, is 2.43 kg. Water in the cylinder adds a maximum of an extra 80 g; some portion of the hose and the fluid inside it contribute to the endpoint weight, depending on the details of the robot configuration, but this is only on the order of a few grams.

The majority of the weight comes from the bearing (350 g) and the structural pieces (the force transducer adds another 320 g). Because of the side loads that the

293
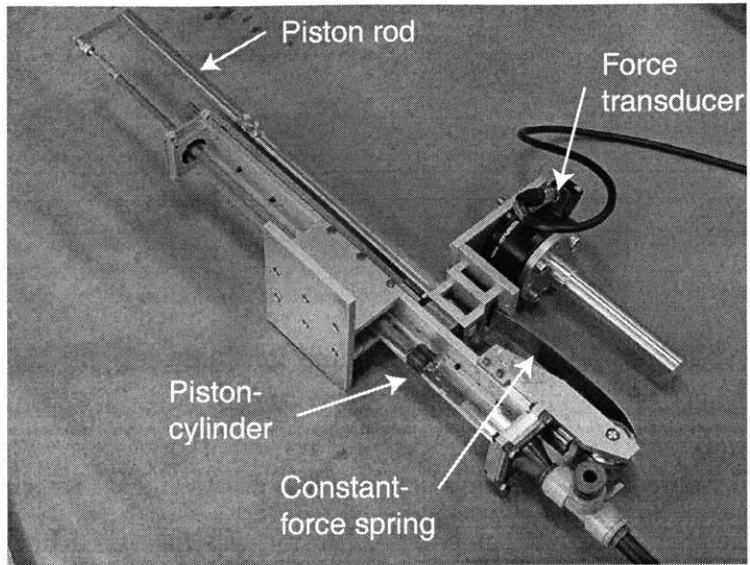
Figure 7-17: Photo of the outer cylinder assembly, the part that mounts to the robot endpoint.
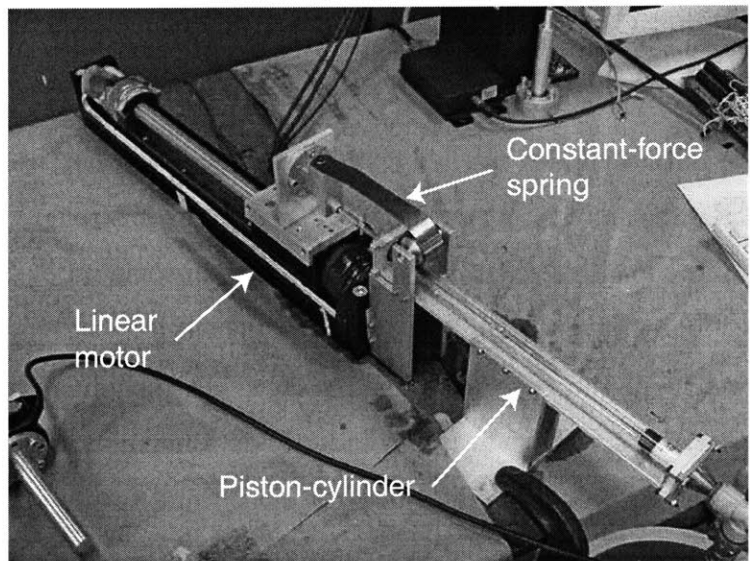
Figure 7-18: Photo of the motor-side cylinder assembly and source motor.

module faces, these are unavoidable; however, because the essential force-generating parts of the module contribute so little weight, these structural members can be smaller than for earlier vertical module designs. Because the electromechanical embodiments weigh around 8 kg, the forces required to support the module's weight ($\sim$ 80 N) are greater than the module output force (65 N). In this fluid prototype, the module weight is just 37% of the output force.

When the key force-generating parts of the fluid module are compared in weight to the force-generating parts of the electromechanical prototypes, a dramatic improvement is evident. The weight of the cylinder, piston, piston rod, and water is approximately 400 g (the vast majority of which is the piston rod assembly). If an additional 230 g accounting for the constant force spring and its mounting assembly are included (as an essential part of bilateral force production, independent of structure), the total is 630 g. By comparison, the total weight of the linear motor forcer and a length of its magnet rod sufficient for 14 inches of travel is approximately 2.72 kg. The total weight of the motor, screw, and nut in the screw-driven module is approximately 2.1 kg. In core force-producing components, the fluid system provides a 77% weight reduction versus the linear motor, and a 70% reduction versus the screw drive.

7.3.2 Force capacity

The linear motor saturates at approximately 70 N of force output in one direction and approximately 80 N in the other direction (the reason for this disparity is unclear). The system was programmed with a high-stiffness (50 kN/m) spring, and the handle on the fluid module was pulled in each direction by hand to its maximum force level. In tension, the system supported almost 60 N, exceeding the force level of the constant force spring (at the expense of lowering the fluid pressure below atmospheric). In compression, the system saturated at around 69 N, due to the saturation of the linear motor. A force trajectory as the system saturates in both directions is shown in figure 7-19. In tension (positive forces), the transmission exceeds its force specification of 45 N, and in compression (negative forces) it exceeds its 65 N specification.
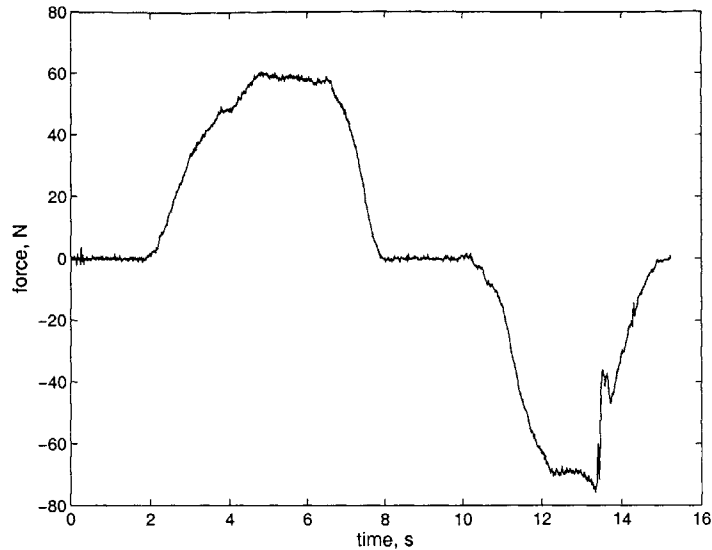
Figure 7-19: Positive and negative force saturation limits of actuator/ fluid transmission system.

7.3.3 Leakage and makeup

The resting leak rate of the system can be measured by allowing the system to sit at rest with the fluid makeup system turned off, and measuring the difference in position at the linear motor and at the output cylinder over time. The two pistons draw closer together as fluid leaks out of the transmission. From the known diameters of the cylinders and the velocity at which the pistons move toward each other, the volumetric leak rate can be determined. Figure 7-20 plots the measured difference between the two piston positions $x_{output} - x_{motor}$ over time. The leak rate is very nearly constant, producing a curve that is very close to linear. The constant relative velocity between the two pistons, as indicated by the slope of a least squares-fit line to the curve in figure 7-20, is 7.45×10^{-4} m/sec (or 268.2 cm/hr). This is between the maximum and minimum predicted values. Assuming equal gaps and leakage at both pistons and working backward with equation 6.29, the gap size can be estimated at 0.0004" radially, or 0.0008" diametrically. This is within the specified range, though closer to the upper bound (of 0.001"). The leakage results in a significant amount of leaked fluid collecting inside the cylinders over time, which must be emptied periodically.

296
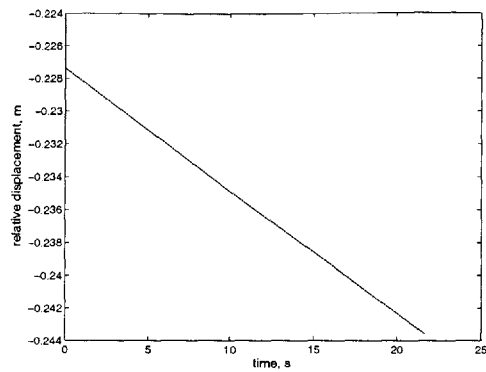
Figure 7-20: Measured difference in motor and output encoder readings versus time, with no fluid makeup. The slope provides the leak rate.
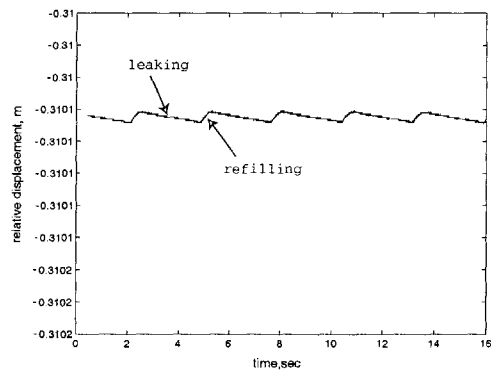
Figure 7-21: Measured difference in motor and output encoder readings versus time, fluid makeup.

A practical clinic-ready version of this actuator would require an additional system (e.g. a suction pump) to remove the excess fluid.

Figure 7-21 shows the effect of the makeup pump. With the system at rest, the periodic refills from the pump keep the relative displacement between the two pistons nearly constant over time. As noted above, the effect of the makeup pump is nearly imperceptible to the touch at the endpoint because the resulting displacement due to the pump's reciprocation is minimal.

297

7.3.4 Virtual stiffness and stiffness measurement

A key objective of the system is to represent virtual springs of various stiffnesses. With simple impedance control at the source actuator, stiffness is implemented. At low frequencies, the endpoint stiffness k is related to the source actuator stiffness k_{src} by the transmission stiffness k_t, with the following series spring relationship:

$$k = \frac{k_{src}k_t}{k_{src} + k_t} \tag{7.7}$$

By measuring the endpoint stiffness and comparing to the motor stiffness, the transmission stiffness can be determined.

With several different values of k_{src}, the system was actuated slowly (with speed less than 0.05 m/sec) by hand. An example of the resulting data is shown in figure 7-22. In this case the source stiffness was set to $k_{src} = 500$ N/m. In time, the plot starts close to the center, positive force is first applied, and the curve is traversed clockwise. The spacing between the upward moving curve and the downward moving curve is due to the Coulomb friction in both directions. The fact that this varies significantly with position suggests position-dependence in the friction. (Of course, because the force is applied by hand, this is another variable). The variation in forces with position is discussed in the next section.

The small peaks that appear occasionally on the curve are the result of a stick-slip phenomenon [5]; they correspond to locations where velocity goes to zero, and local static friction requires additional force (1-3 N) to re-initiate movement. This appears in the plot as a result of the fact that tests were conducted with an imperfect human hand, but it does provide some additional data on the friction in the module. If the hand is moved slightly faster, it is easier to avoid stopping. Figure 7-23 shows another trial where the speed is as high as 0.08 m/sec. Fewer static friction peaks appear in this plot, but the overall friction is slightly higher, because of viscous damping.

For six trials, the force-position curve was fit with a line. The effective transmission stiffness was then computed using equation 7.7. k_{src} was determined by fitting a line to a similar test run at the motor output with the same programmed stiffness. The
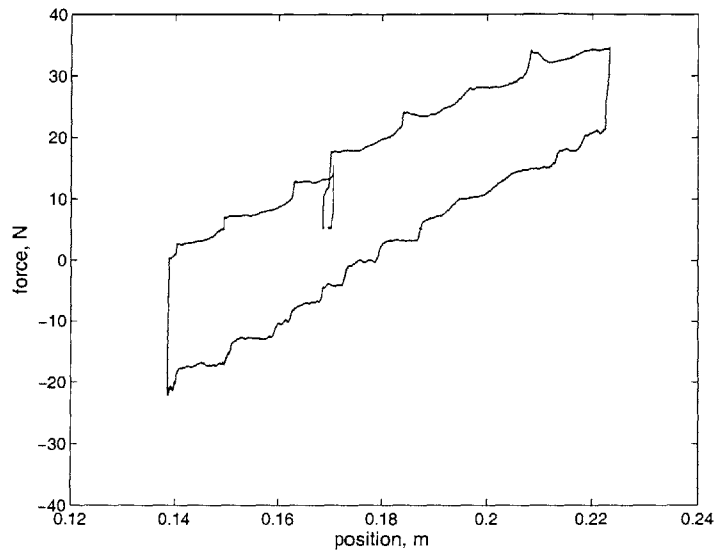
Figure 7-22: Actuation against a 500 N/m virtual spring by hand. Speed always less than 0.05 m/sec.
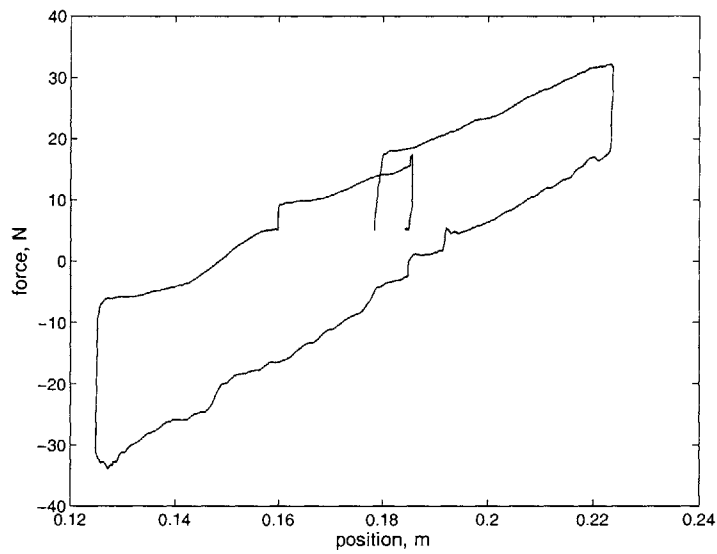
Figure 7-23: Actuation against a 500 N/m virtual spring by hand. Speed as high as 0.08 m/sec.

average of the computed k_t revealed a transmission stiffness k_t of 4020 N/m. This exceeds the specified maximum system stiffness of 2000 N/m, and when coupled with the linear motor's maximum achievable stiffness of 50,000 N/m, permits a maximum system stiffness of at least 3700 N/m. The system's ability to implement a virtual stiffness sufficient for the application is limited only by the endpoint friction, and the variation in that friction force.

7.3.5 Friction measurement

In order to characterize the endpoint friction, the system was driven from the handle at constant velocity, and the force was measured. The linear motor issued nominally zero force for this test, but remained connected to the transmission, such that the measured friction includes the friction of the motor in addition to that of the fluid transmission. The screw-driven module was used as a motion source, with a high-gain PD controller servoing to a constant-velocity reference trajectory at a variety of speeds. A photo of the test setup is shown in figure 7-24. The linear encoder on the output cylinder was used to identify the portion of each test for which the system actually moved at constant velocity (any transients were excluded). The force was measured with the force transducer on the endpoint. The force was averaged over the interval at which constant velocity was maintained, providing a single force value for each velocity. The test was repeated for velocities ranging from 0.001 to 0.1 m/sec, and was done for the module moving upward and downward.

Figure 7-25 shows two of the raw force traces versus position at 0.008 m/sec, one for upward motion and the other for downward motion. Both curves vary substantially with position. The curve for downward travel shows even more variation than that for upward travel, as is typical of this data.

Figure 7-26 shows the average force versus constant velocity for upward travel. Also shown is the predicted friction force, made up of the model-predicted line friction from the major head loss (at these low speeds the minor losses are negligible) as well as 5 Ns/m of linear damping and 3 N of Coulomb friction from the linear motor, as determined in [86]. The model predicts the onset of turbulent flow at 0.07
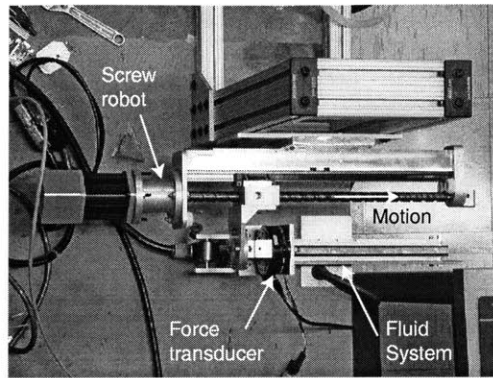
Figure 7-24: Screw robot module used to drive the fluid system endpoint module at constant velocity. Forces measured with force transducer between the two modules. The direction of motion is indicated.
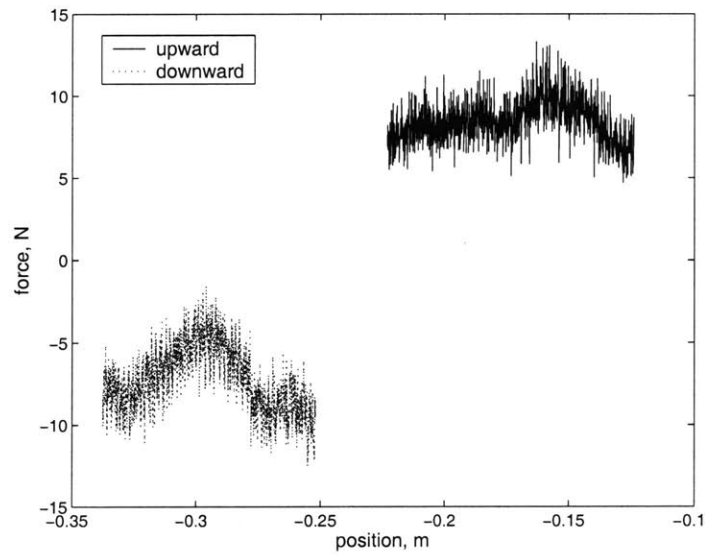
Figure 7-25: Force versus position at a constant speed of 0.008 m/sec, upward and downward.
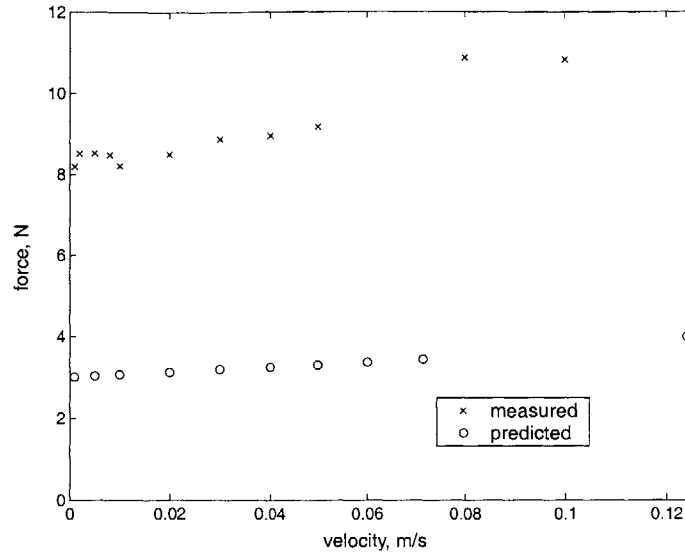
Figure 7-26: Measured and predicted friction versus velocity, upward motion. Prediction includes major fluid damping and Coulomb friction from source motor.

m/sec; below this speed the damping is linear. The total predicted linear damping coefficient for below 0.07 m/sec is 6.24 Ns/m. A linear fit of the measured damping for velocities in this range gives a damping coefficient of 17.1 Ns/m as well as 8.3 N of Coulomb friction. The damping coefficient is significantly higher than predicted; this is undoubtedly partially due to the fact that the hose is not a perfectly straight pipe, as modeled. The linear guide probably contributes additional viscous damping. The more significant concern, however, is the Coulomb friction. The linear guide likely contributes additional unmodeled Coulomb friction, but the total is higher than expected.

Figure 7-27 shows the average friction force versus constant velocity for downward travel. As compared to the upward travel, there is greater irregularity in the force levels; force does not strictly increase monotonically with velocity. This is probably a result of the greater variation of friction with position, as shown in figure 7-25, and the fact that the constant-velocity region covers a slightly different part of the travel for each different velocity. A linear fit of the tested velocities below 0.07 m/sec reveals a linear damping coefficient of 12.0 Ns/m and 8.1 N of Coulomb friction. The
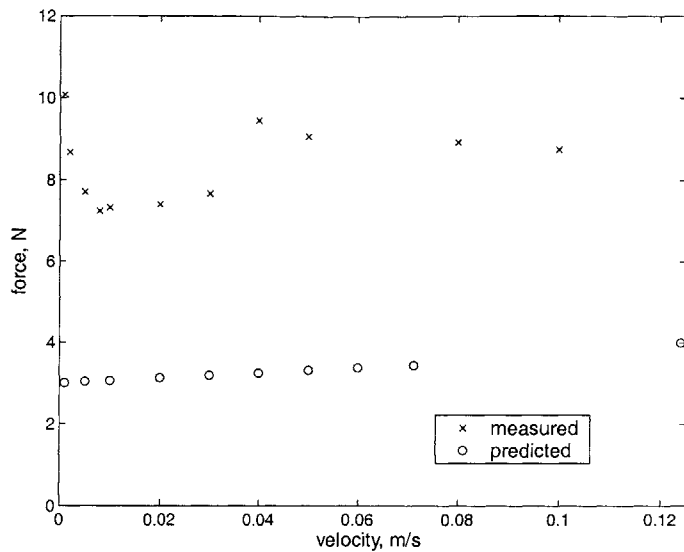
Figure 7-27: Measured and predicted friction versus constant velocity, downward motion. Prediction includes major fluid damping and Coulomb friction from source motor.

Coulomb friction is significantly higher than expected, and its variation with position is especially disappointing.

The virtual stiffness and constant-velocity friction tests both reveal two unmodeled phenomena: the Coulomb friction is higher than expected, and the force to backdrive the system varies considerably with position. The mean friction is around 8 N, and the total force variation is around 7 N. In addition, the virtual stiffness tests exhibited some stick-slip behavior, indicating 1-3 N of additional stiction. Several components and subsystems were tested independently to determine the source of these unwanted problems.

Outer assembly bearing and constant-force spring

The fluid system was disconnected from the piston rod, leaving the constant-force spring's stresses and the friction from the spring, the rotary bearings to which it is mounted, and the linear bearing as the only significant sources of force. The system was driven at a slow but constant velocity (0.008 m/s) with the screw-driven module, and forces were measured. The spring was first slowly extended, then retracted after a

303
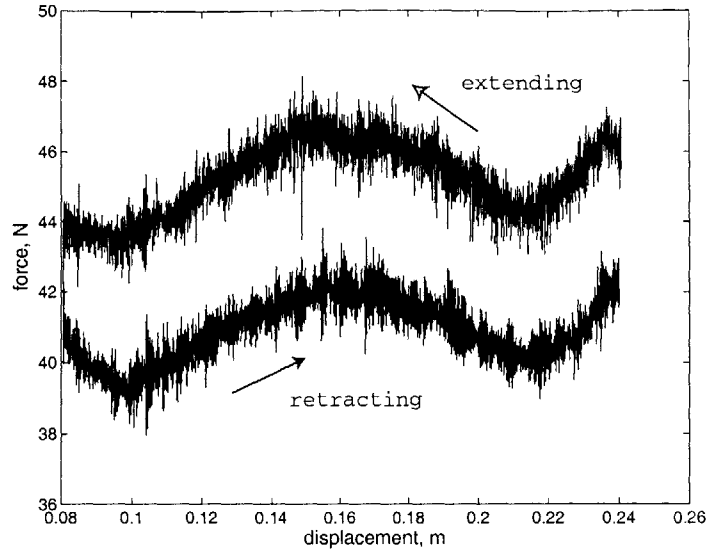
Figure 7-28: Force versus position, outer cylinder constant-force spring and linear ball slide. Tested at a constant speed of 0.008 m/sec.

brief pause. Figure 7-28 shows the force versus position for this test. Several features of this plot are noteworthy. The mean force of 43 N is 9% below the nominal force of 47.1 N. The "extending" and "retracting" portions of the curve are, on average, 4 N apart, indicating an average of at least 2 N of friction in each direction. Finally, there is significant variation in the force level that is repeatable with position. The total variation in a single direction is as much as 3 N. The completely retracted constant-force spring has an outer diameter of approximately 39 mm and therefore a circumference of around 0.12 m. This is consistent with the periodic variation of the force shown in figure 7-28. Possible sources of this periodicity are discussed in the next section, which addresses the other constant-force spring.

Motor assembly constant-force spring

The spring used on the source actuator side of the fluid transmissions was also tested. It was kept in the same mounting as when the actuation system operates, and its end was attached directly to the force transducer on the screw robot. As in the previous section, it was extended and subsequently retracted at constant velocity. In this case,
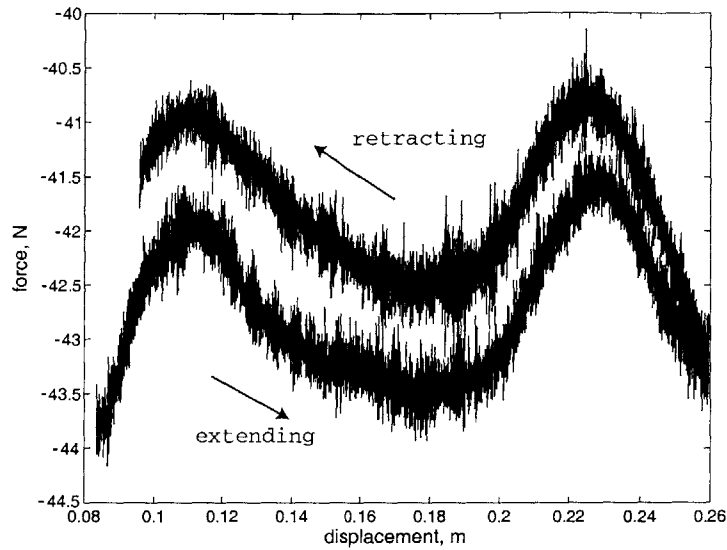
Figure 7-29: Force versus position, constant force spring from source actuator side of transmission. Tested at a constant speed of 0.008 m/sec.

friction results only from the spring and its mounting bearings. The resulting force is plotted versus position in figure 7-29. Again, the average force level (42 N) is lower than the published value of 47.1 N. The difference between the two curves, signifying twice the friction level, is approximately 1 N, suggesting a modest 0.5 N of friction from the constant-force spring and its rotary mount bearings alone.

Like the other spring, the output of this spring varies periodically with its rotation, with a total variation of at least 2 N. It may be possible to eliminate this force variation mechanically either by improving the mounting or by selecting better springs (or both). However, even if this is not feasible, it should be possible to carefully calibrate both springs and use the source actuator to compensate for the difference. The system operates properly as long as the forces from the two springs are equal such that they offset each other. Because the force is highly repeatable with position, and because the position of both springs is sensed constantly with linear encoders, it should not be difficult to add an additional term to the control law that compensates for any instantaneous difference between the two springs by either adding or subtracting force at the actuator side.
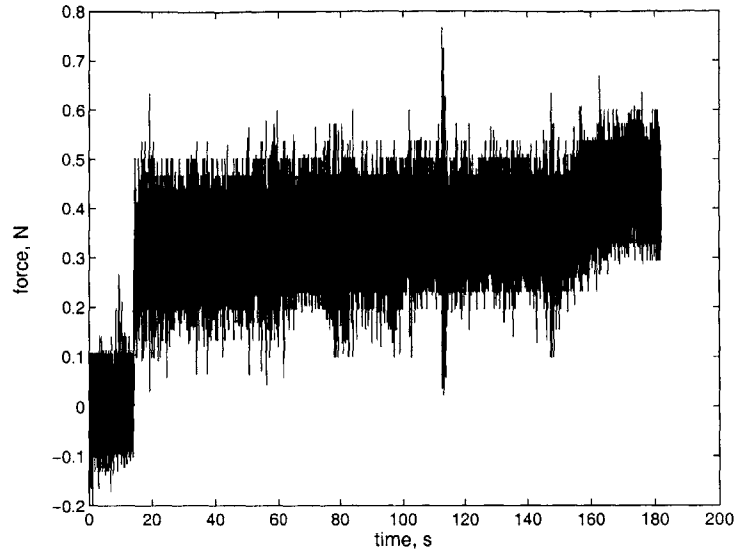
Figure 7-30: Force versus position, pistons and fluid system. Tested at a constant speed of 0.001 m/sec.

Pistons and fluid system

To test the Coulomb friction in the remaining parts of the system, the piston rods on both sides were disconnected from their bearings and mounting hardware (including the constant-force springs). This left only the pistons and the fluid in the system free to move. The screw robot module was used to push the output piston at a constant velocity (0.001 m/sec), and force was measured. In theory this test measures only the friction of the two pistons and the fluid in the system. However, the motor side piston rod had to rest somewhere, so it was allowed to slide against a small plastic fitting that was placed around the inner rim of the cylinder exit. The resulting force is shown versus position in figure 7-30. The force ranges from approximately 0.35 to 0.45 N. The position-dependent increase in the last 30 seconds was typical but not entirely repeatable, indicating some variation from trial to trial.

The friction measured in this test, less than 0.5 N, is quite encouraging. This test includes the core force-producing components of the system, and shows that friction is very low and in line with the cylinder manufacturer's specifications. It suggests that most of the unwanted friction comes from the mechanical parts that surround the

Component(s)	Nominal Friction Force (N)	Total Force Range (N)
Source actuator	3.5	0.5
Outer CF spring and linear guide	2	3
Actuator-side CF spring	0.5	2
Both pistons and fluid	0.4	0.1
Total	**6.4**	**5.6**

Table 7.2: Individual component or subsystem contribution to endpoint Coulomb friction and force variation.

system, not from the core technology, and implies that there is room for improvement in the engineering of the device. In the next section, each source of friction and force variation is considered in terms of its contribution to the total system performance.

Total system friction and force variation

From the tests described above and prior characterization of the linear motor [86], the total endpoint friction of the system and the total variation in apparent friction at the endpoint can be examined, component by component. In this way the major sources of both can be identified and solutions can be proposed.

Table 7.2 summarizes the results of the tests reported in the previous section. The second column provides the nominal or mean Coulomb friction force that each component or subsystem contributes to the endpoint friction. The third column gives the total range of forces that each component or subsystem is responsible at the endpoint. In other words, this gives the total variation in force at the endpoint that appears as variation in friction (whether it actually comes from variable friction, or from another mechanical phenomenon, as in the case of the constant-force springs).

The tests suggest a total of 6.4 N of nominal Coulomb friction from the system components, and a total possible force variation of 5.6 N. For the force to actually vary by 5.6 N, each of the sources of variation would have to tend all the way to their minimum together at one instant, and all the way to their maximum at another instant, which is possible but unlikely. By comparison, the constant-velocity tests on the actual endpoint friction reveal mean Coulomb friction of around 8 N, and a total

variation of as much as 7 N (for instance, from 5 N to 12 N of apparent Coulomb friction in one direction). These totals are slightly larger than the component totals, but this is not entirely surprising, because it is likely that the loading due to applied forces and fluid pressure during operation increases friction and the probability of binding in some components. The linear guide and the piston/cylinder interface may be particularly vulnerable to such loading. An effort has been made to isolate the piston from side loads by including two ball joints in each piston rod assembly, but the ball joints are not frictionless and it is likely that some loading occurs. This could explain the difference between the measured component totals above and the measured total system friction, as well as the additional 1-3 N of static friction described above.

Reducing apparent friction

Table 7.2 offers some insight on how best to improve performance. Of the components, the source actuator is the single largest contributor to the endpoint Coulomb friction. It is inevitable that the actuator friction be reflected at the endpoint in this configuration. It is possible that a redesign of the motor's bearings may offer significant improvement. Alternatively, a force sensor could be placed at the motor to locally reduce the apparent impedance; these results suggest this may be worthwhile. The transmission's compliance and damping are likely to permit more aggressive force feedback than might otherwise be possible, as shown in chapter 5.

Because the friction is so low in the piston/fluid subsystem, it should be possible to improve the bearing arrangement at the output to reduce side loading and therefore the total endpoint friction. One significant source of side loading is the offset between the piston and the bearing, which results in large bending moments applied to the piston rod. A symmetrical arrangement might improve things substantially.

The constant-force springs are the leading source of the apparent variation in endpoint friction. Several possible ways to remedy this problem were suggested in the previous section, including mechanical fixes and calibration.

308

7.3.6 Inertia estimation

The endpoint inertia was estimated by measuring the system's initial condition response against a virtual spring. The linear motor was used to introduce a virtual stiffness. The handle was then displaced and released. The motion trajectory of the handle as it oscillates about the spring's reference position reveals its natural frequency, which can be used with knowledge of the stiffness to estimate the endpoint inertia.

To measure the bulk inertia of the fluid system and linear motor forcer, the compliance of the linear motor controller, rather than the transmission, should dominate the response. Thus the programmed stiffness should be significantly smaller than the fluid transmission stiffness (of approximately 4000 N/m). On the other hand, inertia can be estimated with much greater confidence if the system overshoots; because of moderate frictional losses, the effective damping ratio is too large to permit overshoot when the stiffness is very low. The most promising balance between these two objectives was found with a nominal virtual stiffness of 2000 N/m. Combining in series with the 4000 N/m stiffness of the fluid system, this produces an equivalent endpoint stiffness of approximately 1330 N/m. The initial condition response under this condition is plotted along with a model of a second-order system in figure 7-31. The inertia used in the model shown is 4.3 kg (the damping is 26.5 Ns/m). The fluid system is estimated to have an apparent endpoint inertia of 2.52 kg. The linear motor forcer inertia is 1.45 kg. The moving hardware at the output cylinder has mass measured at 0.93 kg, giving a total predicted endpoint inertia of 4.9 kg. Thus the estimated and predicted inertia agree within 12%. Because the model assumes a lumped mass at the end of a massless spring, but the real system inertia is distributed along part of the spring (the 4000 N/m part due to the fluid transmission compliance), it is not surprising that the measured inertia is less than the predicted inertia.
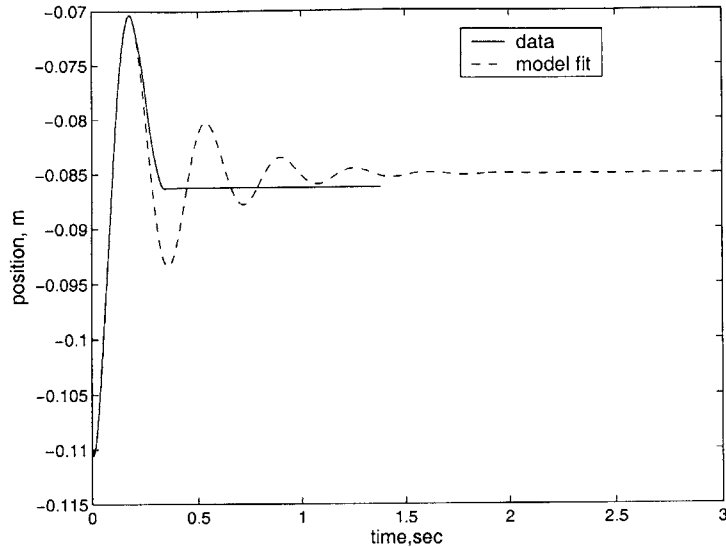
Figure 7-31: Initial condition response of the fluid system, along with a model fit. Model parameters are 1333 N/m stiffness, 26.5 Ns/m damping, and 4.3 kg inertia.

7.3.7 High-force bandwidth

Although not a core requirement for this application, it is useful to evaluate the bandwidth of the actuation system. For force-controlled actuators, a common measure is the high-force bandwidth [53, 83], which quantifies the response of the output force to an input force command. To evaluate the force bandwidth for this actuation system, the output was locked by clamping the handle in a bench vise. Sinusoidal forces were commanded at the linear motor with an amplitude of 12 N, or approximately 20% of the peak force specification. The output force was measured by the force transducer to which the handle is mounted.

The magnitude response is plotted versus frequency in figure 7-32. The system has a resonant peak around 8 Hz, with a damping ratio between 0.1 and 0.2. The open-loop bandwidth, where the magnitude crosses -3 dB, is approximately 20 Hz. Considering that voluntary human motion is generally limited to less than 5 Hz, this is sufficient for the application.
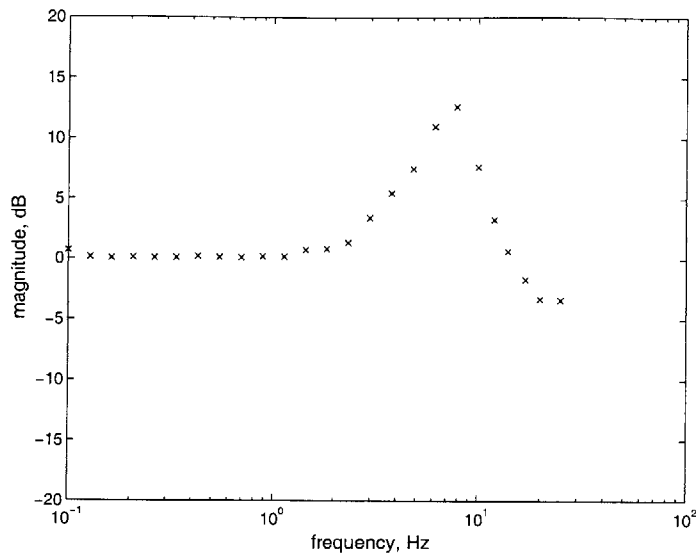
Figure 7-32: Output force magnitude over input force magnitude.

7.3.8 Conclusions: Limiting factors and suggestions

The fluid transmission prototype presented in this chapter demonstrates the concept of dramatically reducing effective actuator weight while (mostly) preserving the low-impedance benefits of electromagnetic actuation. The prototype endpoint package is 70% lighter than each of two electromechanical prototypes. The core force-producing components of the system, the piston-cylinders, the fluid, and the constant-force springs, are lightweight and have low mechanical impedance. These positive results show future promise in this approach.

Despite these accomplishments, the present design is limited by several factors. Coulomb friction remains a factor of four higher than specification (the source actuator itself has almost twice the target friction). Some of the reasons for this and ways to reduce the friction are discussed in the preceding sections. Similarly, the reflected inertia is higher than specified. This is mostly a result of the fluid line geometry, and can only be significantly reduced by increasing the line diameter or reducing its length (particularly the former). This is a tradeoff that must be managed at the design phase. It might even prove beneficial to make a composite line that is very large in some regions and narrow in regions where it must be flexible or lightweight.

311

While the friction and inertia are higher than desired, the leakage rate is substantial. Although the fluid makeup system adequately replaces lost fluid, the spillage is significant. All known methods of reducing leakage tend to increase endpoint impedance, either directly via increased seal friction or indirectly via increased reflected inertia and viscous drag (if the operating pressure is reduced; see chapter 6). In order to keep friction low, it is likely that the leak rate cannot be reduced by the order(s) of magnitude necessary to make it negligible. Instead, an additional system should be incorporated to remove the leaked fluid. Suction is a possibility; the suction system need not operate at very high pressures nor handle large flow volume.

Finally, the simple makeup system used here is vulnerable to drift as forces are applied to the transmission. This can be solved by incorporating a flow-controlled pump that is controlled by the system software. Such a scheme could use the motion and force sensors to monitor leakage and modulate the makeup rate accordingly. The cost is greater complexity and expense.

When high forces, low impedance, and low weight are sought in an actuator, the designer is left with few options. The configuration proposed here, in spite of its flaws, has the potential to solve this problem for certain applications.

Dynamics in fluid transmissions

As shown in chapter 5, it may be advantageous to include specific dynamics in series between a source actuator and an environment. Parallel spring-dampers were shown to be particularly useful, and inertia was shown to have some stabilizing properties as well. The fluid transmission prototype described in this chapter exhibits compliance and damping as well as inertia, and may enable more accurate rendering of endpoint impedance when combined with force feedback, in addition to its benefits in reducing the weight of the actuator package.

If more closely controlled series dynamic properties are desired, the design could be altered to include additional dynamics in an integral or separate module. Appendix B provides analysis and a suggestion for one such design.

Chapter 8

Conclusions, Future Work, and Alternatives

8.1 Combining hardware and control design

Despite the advances in both actuator and control design for interaction described throughout this thesis, there remains no single "silver bullet" to solve the human-robot high-force interaction problem. Using knowledge of the environment can permit more aggressive control that improves performance considerably, but if the physical system impedance is not within an order of magnitude of the target behavior, control alone is unlikely to bring system performance to the desired level. Actuators that are completely non-back-drivable are unlikely to be effective, regardless of control strategy.

Similarly, off-the-shelf actuator technology cannot provide the combination of force density and low impedance to intrinsically solve the problem. Perhaps the best the designer can do with off-the-shelf technology is to design hardware that meets force requirements and to strike a middle ground between impedance requirements and low mass. Force-feedback control can be used to reduce the impedance error, and multiple-degree-of-freedom force feedback may allow some reduction in the apparent actuator mass that adds inertia to the non-coincident robot degrees-of-freedom. This can be particularly effective if knowledge of the environment is incorporated in controller

313

design. Small-reduction gearing can help to reduce actuator mass partially without increasing reflected inertia or friction too dramatically. This approach can suffice, but it can still result in relatively high intrinsic inertia in the robot, which can be dangerous in the event of controller failure.

An alternative approach that holds the promise of more closely meeting all specifications in hardware alone is to design custom actuators. One such approach is described in chapters 6 and 7; several others are described in section 8.3. As yet, none of these methods have solved the problems convincingly enough to dominate the field. However, all have potential for certain applications, and may need refinement to apply more broadly.

8.2 Progress made and future work

This thesis has addressed a challenging problem, the development of high force-density actuators with low mechanical impedance that can stably interact with a range of environments, with a multi-pronged approach.

The inclusion of partial knowledge of the environment in the design of interaction controllers was explored. It was shown that classical control design tools cannot be directly applied to such systems to meet stability and performance objectives, even if the environment is perfectly known, though the tools remain useful for analysis. Robust control tools were applied to enforce a new stability constraint termed *complementary stability*, which was specially formulated to ensure stable interaction with environments with bounded static and dynamic properties. Because linear design tools cannot be directly applied, a numeric parameter search method was proposed that uses the new stability measure and evaluates performance on the basis of robot port impedance. This design approach was applied to a simplified linear robot model motivated by the literature and was shown to converge to a non-obvious control design that improved performance versus state of the art passive controllers. The controller was then implemented on a robot with substantial nonlinear friction, and was shown to improve coupled stability properties as well as performance. The model produces

conservative results that guarantee stability, but are less conservative than passive controllers.

As control alone cannot solve the high-force, low-impedance interaction problem in all cases, changes to the physical actuator structure were considered. The design algorithm developed for controller design was applied to search for optimal parameters for physical "impedance shapers" to improve actuator structure with respect to interactive performance. This search method was applied to a structure proposed in the literature, and provided a systematic way to select parameters in addition to insights about the differences between previously proposed approaches. It was demonstrated that both approaches have similar results despite their apparent difference (the presence or absence of damping) because of the properties of the environment; the differences are reduced when complementary stability is used in lieu of passivity. The design method was then applied to a new impedance shaper structure in the form of a mass in series with the actuator. It was shown that adding mass can permit more aggressive force feedback and improve performance.

Finally, a new design for remote transmissions was proposed. This design leverages the low impedance of an electromagnetic source actuator and the high force density of hydraulic systems. The proposed configuration was modeled and critical tradeoffs between geometric parameters were identified. A prototype was constructed that reduced the effective actuator weight by 70% when compared with two electromechanical designs. Coulomb friction was higher than desired and expected, but it was shown that this resulted primarily from the mechanical elements surrounding the system, not the core technology that makes the design novel. Excessive endpoint inertia and fluid leakage were also problematic. The inertia can be reduced at the expense of other design parameters, as dictated by the model tradeoffs. Reducing leakage is likely to increase endpoint impedance, so it is perhaps better to design a system to manage the leaked fluid.

Progress was made in control, actuator, and transmission technology. It appears that the best solution to this challenge with current technology requires effort on all facets of the problem.

315

8.2.1 Future work on control

Incorporating limited knowledge of the environment was an important step in reducing controller conservatism. This was accomplished here, but this approach requires that the controller structure be provided. An even more useful design method should also provide the desired structure. This may be possible with modern control techniques, using an architecture similar to that used in chapters 3 and 4.

This approach should also be extended to multiple degrees-of-freedom. The single degree-of-freedom case is useful, particularly for actuator design. Much of the unwanted high-impedance dynamics in robots comes from the drivetrain, so closing a force feedback loop around the actuator and transmission elements like gears, screws, etc. can be helpful. However, it may also be useful to design a controller at the robot level, so that the joint actuators can improve endpoint behavior based on forces sensed at the endpoint. The approach detailed here requires a model (though perhaps not an exact one) of the robot's low-order resonant behavior. This may be more difficult to obtain in multiple degrees-of-freedom, particularly as the modes may change significantly with robot pose. To tackle this problem, it may be wise to start with an analysis of the sensitivity of the control technique and the model to changes in the parameters that govern its resonance.

As this approach is extended to a greater variety of systems, it will be important to know how accurate the model must be to develop a reliable controller design. The sensitivity of the control algorithm results to changes in model parameters should be explored and quantified. Particular attention should be paid to parameters that are difficult to estimate, such as the structural stiffness and damping. There is already some indication, from the results of chapter 4, that an inexact model may permit useful control design. However, this should be quantified in order to guide future model-based control design.

The experiments of chapter 4 revealed several interesting differences between the theory of natural admittance control and practice. These are worth exploring, as the controller shows a great deal of promise. Its use of both motion and force feedback

to reduce impedance potentially gives it greater flexibility than controllers that use only force.

8.2.2 Future work on impedance shapers

Chapter 5 provides a design approach for impedance shapers. This approach should be applied to different dynamic structures to search for additional promising configurations. This area is relatively unexplored, with only two or three groups conducting research, and the results for all have been fairly promising. Because of the limits of paper-and-pencil analysis, only extremely simple and low-order dynamic structures have been proposed in prior work. The method provided in chapter 5 offers the flexibility to analyze more complex systems in a straightforward way to see if they are advantageous. The method exports the difficulty to a computer, leaving only simple calculations and the interpretation of results to the user.

Because it is so clear that impedance shapers can help interactive performance, some attention should be focused on their detailed design and packaging. One principle problem is that they generally require additional hardware to create the desired dynamics. Incorporating the desired dynamics into machine elements that must be part of the system anyway (e.g. motor shafts, robot links, etc.) or developing other low-profile implementations could make these much more appealing.

8.2.3 Future work on passive fluid transmissions

The prototype transmission system described in chapter 7 has several flaws that should be corrected. Although the prototype demonstrates the basic concept and substantially reduces effective weight, the next generation could likely reduce friction via a more symmetrical endpoint package design, and handle leakage with a separate "cleanup" subsystem. Alternative bearing designs may be able to reduce both friction and weight.

Fluid transmission technology may open the door to solving some of the more challenging human-interactive robot tasks. These include machines that must bear

part of the human body weight [98] or those that actuate the fingers [67]. The potential of this actuation methodology could be more fully demonstrated through its implementation in a complex, practical device.

Fluid transmissions also provide a platform to implement series dynamics, or impedance shapers. This could permit the use of a geared source actuator, possibly reducing its weight and cost, and improving portability. Appendix B discusses one method for introducing and controlling desired dynamics. Section 8.3.4 below suggests a way of actively controlling series dynamics in fluid transmissions.

8.3 Promising alternative approaches

In this section, several alternative approaches with promise to contribute to the solution of the high force density, low impedance, stable actuation problem are discussed.

8.3.1 Pressure-controlled hydraulics with force feedback

Present embodiments of pressure-regulated hydraulic systems use a servo-hydraulic structure that produces high endpoint impedance, as shown in chapter 2. This impedance can be partially reduced if large leakage flow is allowed, but this comes at the expense of efficiency and compressor size. Excessive leakage flow is also likely to produce noise in the output.

Pressure or force feedback is another way to reduce apparent impedance, but this is likely to sacrifice passivity and compromise coupled stability. However, there is one structural difference between high-impedance pressure control valves and highly geared mechanical actuators that warrants further study of force feedback in hydraulic systems. Adding gearing to an electromechanical actuator increases reflected inertia, friction, and damping, as shown in chapter 2. On the other hand, a servo-hydraulic actuator derives most of its high impedance from the viscous damping that results from forcing fluid through small orifices. Fluid inertia originates from the need to accelerate large amounts of fluid through small passageways of finite length; small orifices introduce substantial head loss, but generally do not require increased accel-

eration of significant quantities of fluid. The energy loss translates into substantial damping, but theoretically little energy storage in the form of inertia. As discussed in chapter 2, in a mechanical force-feedback system, inertia rejection causes passivity to be lost, not friction rejection. It is conceivable that force feedback could be used to reject the increased energy loss due to the orifices, without loss of passivity. This warrants further analysis. Given an appropriate model of a high-impedance pressure-regulated system, the tools derived in this thesis could easily be applied to study this situation.

The passive hydraulic transmission (with electromechanical source) described in this thesis is an alternative method of regulating hydraulic pressure. It circumvents the use of separate servovalves and compressors. Its limitations are rooted in the tradeoffs between low friction and low fluid leakage, and are characteristic of any hydraulic system. Aside from these issues, this approach excels at generating force and impedance, and this is accomplished with extremely simple and robust control. If hydraulics are to be used for human interaction (and there are good reasons to do so), it is difficult to see why a more conventional arrangement would be a better choice. An improvement that eliminates one or more of the major limiting factors in this design, as cited in chapter 7, would make this approach more appealing.

Should a proper hydraulic control strategy produce an effective impedance source without compromising coupled stability, the system would still be subject to many of the limitations found in the study of the hydraulic system in chapters 6 and 7 The system would be vulnerable to leakage and seal friction. It would require a low-friction bearing. If the servovalve is located remotely from the end-effector, the fluid line inertia and damping are apparent at the endpoint, and there is no obvious advantage as compared to the architecture suggested here. If the servovalve is located at the end-effector, then the line impedance is upstream of the actuator and sensors, and could be eliminated through proper control. The tradeoff in this arrangement is that the servovalve must be located at the actuator, and its mass contributes to the total actuator mass and weight. This approach seems to be the most promising for conventional pressure hydraulics. It depends on two enabling technologies: a

319
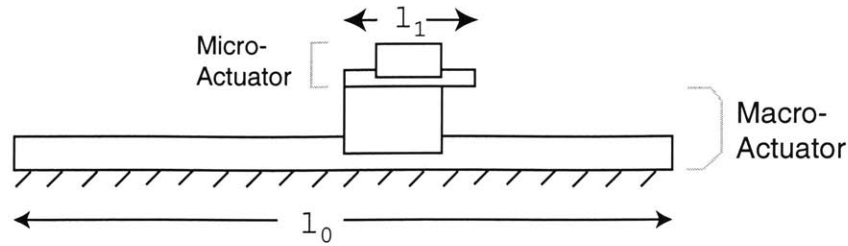
Figure 8-1: Schematic representation of a macro-micro linear robot.

low-mass pressure-regulating servovalve, and pressure- or force-feedback control that rejects the head loss introduced by the servovalve without compromising coupled stability. These two issues should be studied, and if promising results are found, this design should be explored further.

8.3.2   Multiple actuators in series

Because a single actuator is not readily available to satisfy the force, impedance, mass, and range-of-motion requirements, solutions that use multiple actuators for each degree of freedom are considered. One way to reduce the weight of an actuator that meets the force and impedance requirements is to reduce its travel; linear electromagnetic actuators, for example, can produce high forces with extremely low impedance and mass when their travel is short.

A configuration that uses a short-stroke actuator to improve the fidelity of the force output along with a long-stroke, high-impedance actuator is shown schematically in figure 8-1. In this configuration the two actuators are mounted in series, with the "macro" actuator mounted to ground (or a robot link), and the "micro" actuator mounted to the macro actuator. Because they are in series, both actuators must satisfy the system force requirements. The total workspace is defined by the travel of the macro actuator, l_0, plus the travel of the micro actuator, l_1.

In one extreme case, if l_1 is large enough to accommodate all endpoint motion relative to the virtual trajectory, then the macro actuator does not need to move in response to endpoint forces or motion. In this case, it can act solely as a motion

source, and any motion controller (that does not provoke coupled stability problems) is acceptable for this actuator; it is okay for the impedance to be large. The micro actuator can then provide all of the desired endpoint dynamics via its controller. Assuming that it acts as a quality force source, this should be quite effective. This is similar to the extreme solution for spring-damper impedance shapers described in chapter 5, with the series elements providing all of the target dynamic behavior. Here, these series dynamics could be directly regulated, forming a sort of active impedance shaper.

Depending on the application, it may be self-defeating for l_1 to be so large. If relatively large deviation is expected from the virtual trajectory, the travel required of the micro actuator is likely to be large enough to increase the mass beyond the desired value. In this case, l_1 must be less than the expected displacement from the virtual trajectory, and the macro actuator must move in response to the position or force of the endpoint, to prevent the micro actuator from reaching the limits of its travel. This case is analyzed below.

The desired endpoint behavior is often that of a virtual spring, grounded at the virtual trajectory. If the controllers for each actuator are used to implement virtual springs, the result can be modeled as in figure 8-2. The macro actuator has high impedance, modeled as the inertia M_0 and the friction force F_f. Its motion is defined by x_0, and its controller creates a virtual spring of stiffness k_0 between the actuator endpoint position and ground. The micro actuator has much smaller mass $M_1 << M_0$, and is assumed frictionless. The micro actuator controller creates a virtual spring k_1 between the endpoint positions of the two actuators.

If the system has a target endpoint stiffness k_{des} and a maximum force F_{max}, which is incurred in response to deviation of the endpoint position from the virtual trajectory and the linear spring k_{des}, the maximum travel of the actuator in response to applied endpoint forces is $\frac{F_{max}}{k_{des}}$. If $\frac{F_{max}}{k_{des}} > \frac{l_1}{2}$, the travel of the micro actuator is insufficient. In this case, some movement of M_0 is required to keep M_1 from hitting its stops. The desired endpoint stiffness can be achieved (at low frequencies) via a
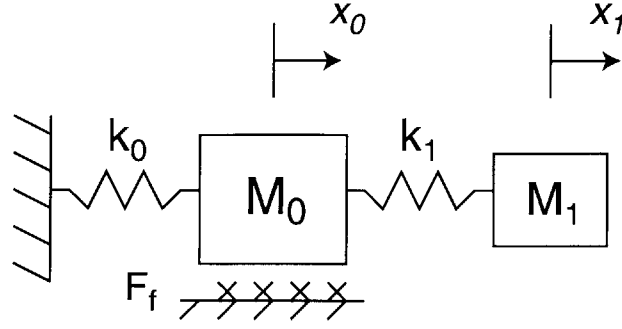
Figure 8-2: Model of an impedance-controlled macro-micro robot.

series combination of the springs k_0 and k_1, such that:

$$\frac{1}{k_{des}} = \frac{1}{k_0} + \frac{1}{k_1} \tag{8.1}$$

The minimum stiffness k_1 such that M_1 does not reach its stroke limits is $k_1 = 2\frac{F_{max}}{l_1}$. Using this value and equation 8.1, the stiffness of the macro-actuator can be determined:

$$k_0 = \frac{2k_{des}F_{max}}{(2F_{max} - l_1 k_{des})} \tag{8.2}$$

The total apparent endpoint inertia (for frequencies below the resonance $\sqrt{\frac{k_1}{M_1}}$ is:

$$M_{app} = M_1 + M_{0,app} \tag{8.3}$$

where $M_{0,app}$ is determined by the ratio of $\ddot{x}_0$ to $\ddot{x}_1$. Statically, under some applied force F:

$$
\begin{aligned}
x_0 &= \frac{F}{k_0} \\
x_1 &= \frac{F}{k_{des}}
\end{aligned} \tag{8.4}
$$

Using 8.2 and 8.4 and differentiating twice produces:

$$\ddot{x}_0 = \frac{(2F_{max} - l_1 k_{des})}{2F_{max}} \ddot{x}_1 \tag{8.5}$$

When $l_1 k_{des} \geq 2F_{max}$, the micro actuator provides sufficient travel, the macro actuator

322

does not need to move in response to endpoint forces, and the extreme case described above is realized. When $2F_{max} > l_1 k_{des}$, motion x_0 is required, but is somewhat less than the motion x_1. The apparent mass is:

$$M_{app} = M_1 + \frac{(2F_{max} - l_1 k_{des})}{2F_{max}} \tag{8.6}$$

A second extreme case is evident from equations 8.5 and 8.6. If the micro actuator travel approaches zero, the inertia of the macro actuator dominates, and its impedance properties become increasingly relevant.

Plugging in some numbers relevant to the vertical robot application considered throughout this thesis provides a check on which case is most applicable. The vertical robot application requires forces on the order of 50 N; a typical stiffness used in robotic therapy is 200 N/m. A high-force voice coil actuator from Servo Magnetics, Inc. can produce 9.5 lb (42 N) force continuously (25 lb peak) over .018 in (.005 m) of travel, in a package that weighs only 20 oz (0.5 kg), with a 4 oz inertia for the moving part. Plugging $F_{max} = 50, k_{des} = 200$, and $l_1 = 0.005$ into equation 8.6 produces:

$$M_{app} = M_1 + 0.99 M_0 \tag{8.7}$$

For these parameters, 99% of the inertia of the macro actuator, which is presumably much larger than desired, appears at the endpoint. To prevent the micro actuator from hitting its stops, k_1 must be made so stiff that almost all of the motion comes from k_0. The only way to reduce $M_{0,app}$ is by increasing either the stiffness k_{des} (which is specified by the problem requirements) or the travel l_1 (which inevitably increases endpoint mass and inertia). Furthermore, if the application requires lower stiffness with the same maximum force, the problem gets worse.

It is likely that a nonlinear control scheme could be devised that permits motion of only the micro actuator when travel is short, and drives motion of M_0 when M_1 approaches its travel limits. The core problem remains, however, that the macro actuator must move in response to applied endpoint forces when the travel is large. If the apparent inertia of the macro actuator is reduced with control, the system becomes

vulnerable to coupled stability problems, as discussed in chapters 2, 3, and 4. This is true whether M_0 moves in response to measured force or measured motion x_1. Coupled stability problems may be slightly ameliorated because the micro actuator provides soft dynamics between the macro actuator and the environment. In this sense, the micro actuator fills the same role as the passive mechanical filters described elsewhere in this thesis. The main advantage of a series actuator rather than a passive series mechanical filter is that the actuator can be tuned to provide different dynamics at different times. The cost is the greater mass, greater expense, and greater complexity of a second actuator.

This analysis assumes a perfect force source for the micro actuator, and holds true regardless of its implementation. This approach appears to have especially high value when the travel relative to the virtual trajectory is small compared to the total robot workspace. This may be true for robots used in assembly, deburring, grinding, or similar tasks, where the total workspace might be large but the location of parts is known within a few millimeters. For the human robotics problem, however, peak forces are generally reached when the subject reaches a maximum excursion from some desired position on the virtual trajectory. In general, the robot must maintain the flexibility to apply forces throughout most or all of its workspace, regardless of the location of the virtual trajectory. This constraint reduces the potential benefits of a macro-micro actuation approach.

8.3.3 Multiple actuators in parallel

Along the same lines as the macro-micro approach described above, Zinn et al. have proposed using two actuators in parallel for a single degree-of-freedom [124]. In this approach, there is a "macro" actuator at the base, connected to the endpoint through a remote transmission such as a cable drive. This actuator is responsible for large forces at low frequencies. This is augmented with a second, smaller actuator located at the endpoint. The smaller actuator provides the desired high-frequency content, but not the full force capacity.

While the small motor is likely to provide desired high-frequency behavior effec-

tively, it may not be able to reject unwanted vibrations from the cable as successfully. To do this might require an elaborate estimation and control scheme. An improvement might use a hydraulic transmission like that proposed in chapters 6 and 7 for the low-frequency, high-force behavior, and a small electromagnetic actuator for high-fidelity, high-frequency performance. It would be necessary to make the friction in the macro system as low as possible, as it appears directly as part of the endpoint impedance, but it would alleviate any high-frequency demands from the fluid system. This might produce substantially improved haptic immersion.

8.3.4 Actively regulated series dynamics

The fluid system proposed in chapters 6 and 7 may permit the direct actuation of series dynamics, so that they can be made time-varying. If a third piston-cylinder set is added to the transmission, an additional actuator can be used to control the dynamics of this piston. If this actuator is locked in position, the design is approximately equivalent dynamically to the transmission detailed in chapter 7. If the actuator is allowed to move with some specified dynamic behavior (e.g. stiffness and damping), these dynamics act as a series impedance shaper. (This may be more clear from appendix B). This could permit online, active variation of series dynamics, and may have potential for partially decoupling the power and force production functions of the source actuator from its representation of dynamic behavior. Comparing the performance of such a scheme (in theory and practice) to other architectures that use redundant actuators for this application might present an interesting research opportunity.

There are at least three means of implementing series dynamics in some form. Impedance-shaping series elements can be fixed and purely passive, as assumed in chapter 5. Several possible means of implementing these shapers in a hydraulic system are discussed in appendix B. Alternatively, the series dynamics could be separately actuated by an independent actuation element that connects remotely to the transmission. This is the configuration described in this section. Finally, a series actuator could be used to fully and directly actuate series behavior, as described in section

325

8.3.2. The analytical tools developed in chapters 3 through 5 may be adaptable to explore these options and predict which is likely to be most fruitful.

8.3.5 Direct impedance modulation

A complementary approach to implementing interaction control without relying on feedback, and the threat to stability that it brings, is to modulate impedance directly. This is similar to the idea of impedance shaping described in chapter 5, but requires that the physical impedance be variable so as to introduce different endpoint impedances. This approach is used by humans for manipulation; by the use of redundant actuators (muscles) for each degree-of-freedom, the output impedance can be changed without changing the endpoint force, for example via co-contraction [58, 17].

Although this approach has not become prevalent in robotics, it can be implemented in a number of ways. Field-responsive fluids such as electro-rheologic or magneto-rheologic fluids can provide variable damping when the applied electrical or magnetic fields is changed. An electromechanical actuator design that provides for direct modulation of impedance has been proposed [41].

Artificial muscle technology, for example electro-active polymer muscles, may provide improved opportunities to implement actuation schemes that permit the modulation of impedance without feedback. This field is maturing and holds promise for the future of machines that can act and interact like humans.

8.3.6 Monopropellant actuators

Goldfarb is developing actuators that use a monopropellant fuel (such as hydrogen peroxide) and a chemical reaction to produce pressurized vapor [48, 43]. By controlling the amount of fuel that is reacted, output power can be regulated. This approach avoids the weight of batteries and electromagnetics, and looks especially promising for mobile robots. By using a compressible fluid, this approach is vulnerable to many of the drawbacks of pneumatics that result from energy storage in the fluids. These problems may be manageable however, and may be preferable to the drawbacks of

326

other approaches.

8.4 Future direction for improving interaction

Although the problem of low-impedance, high-force stable interaction is an unsolved challenge, the set of design options is somewhat limited. Because there are no "silver bullets" available, new technologies should be evaluated on their merits for solving this problem as they are developed. Designers should keep an open mind to evolving technology.

Presently, there are few technologies that are capable of the appropriate force densities for the applications discussed throughout this thesis that also are equipped with the ability to render a range of impedances accurately.

Work in this field should, for the time being, focus on actuation and power transmission methods that can provide sufficient force densities (such as hydraulics, pneumatics, tension elements such as cables, etc.). The most promising innovation lies in finding new ways to make these technologies render impedance more effectively, either through new design configurations or control. This thesis is an example of this approach. Conventional actuation approaches have been used but have been configured in a simple yet novel way to help them render impedance more effectively, while more of the information available has been leveraged to reduce conservatism in control. The specific approaches described here have shown promise but similar, system-level design changes should also be examined as this field of study continues to evolve.

Appendix A

High-frequency impedance phase behavior

In figures 4-12 and 4-14, the phase of the port impedance under lag and lead control, respectively, tends to a high positive value at frequencies beyond the structural resonance. This can be seen more clearly in figure A-1, which expands the axes to show that the port impedance phase tends toward 2π at high frequencies, rather than to zero. In this respect the impedance is fundamentally different from the uncontrolled system, and from the system with lower-gain force feedback controllers. This is a true characteristic of the port impedance phase, not an artifact of phase wrapping. This characteristic is not present when the force feedback gain is low, but is observed with proportional, lag, and lead control when the gain is raised sufficiently. This appendix explores the source of this characteristic and considers its implications for robot control and the application of the design method proposed in chapters 3 and 4.

A.1 Port functions and nonminimum phase

The change in the final value of the phase as the force feedback gain is increased is significant. Because the number of poles and zeros in the impedance function remains constant regardless of (nonzero) feedback gain, the relative order and therefore the phase at high frequencies should remain constant as long as the system is minimum
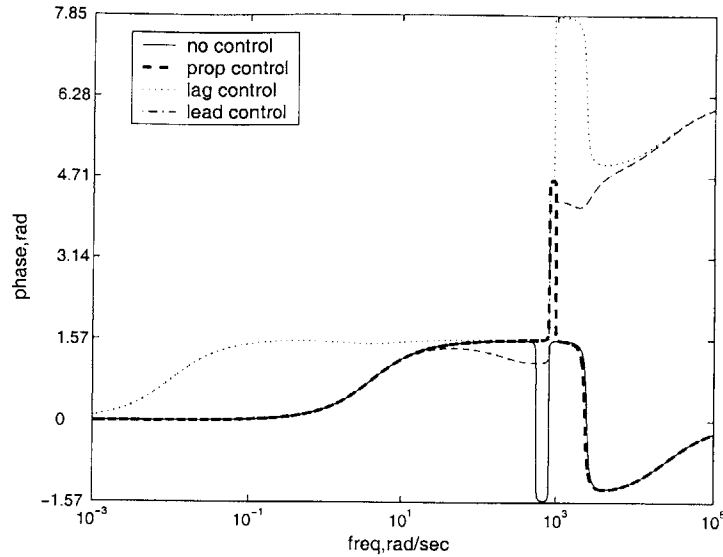
Figure A-1: Port impedance phase for single-resonance system with proportional, lag, and lead force feedback, and without control.

phase. In fact, the change in high-frequency phase coincides with the movement of a pair of complex poles of the port impedance function from the left half to the right half of the complex plane (the real value of the poles changes from negative to positive). With two right-half-plane poles, the function represents a nonminimum phase system. Each pole in the right-half plane contributes $+\frac{\pi}{2}$ radians of phase (rather than $-\frac{\pi}{2}$ if in the left-half plane) to the frequency response, producing the behavior shown in figure A-1.

A.1.1 Right-half-plane impedance poles and coupled stability

Poles in the right-half plane are generally indicative of a loss of stability, and this case is no exception. Because a port function formulation is used, however, care must be taken in understanding the conditions under which instability is predicted. Figure 3-1 shows the connection of two port functions, an impedance Z and an admittance Y. Note that when $Y = 0$, the loop is broken and "closed-loop" behavior is equivalent to the "open-loop" behavior defined by the port function Z. In this case, if Z has one or

330

more unstable poles, the system is unstable. If Z represents a robot impedance and Y an environment admittance, then the condition $Y = 0$ represents an environment with zero admittance, or infinite impedance. Such an environment is a kinematic constraint. Thus the "unstable" robot impedance poles predict that the robot will not couple stably to a kinematic constraint. Because the class of environments used to design the controllers in chapters 3 and 4 does not include rigid constraints, controllers that place impedance poles in the right-half plane do not necessarily compromise complementary stability (though they ensure that the port is nonpassive).

A.2 Port impedance without force transducer

Further insight can be gained into the predicted physical implications of the right-half-plane poles by considering the port impedance on each side of the force transducer. Figure A-2 shows the phase of the port impedance of the robot system under proportional force feedback with a gain of 5. Two traces are shown, representing the phase of the impedance at two different physical ports indicated by "A" and "B" in figure A-3. (In both cases ideal force feedback is presumed; the applied force at port A is always equal to the force at port B, but the motion at the two ports, and hence the impedance, differs). The phase at port A (the location used throughout chapters 3 and 4) shows behavior characteristic of a nonminimum phase system. The phase at port B, however, is well-behaved; all poles and zeros of this impedance function are in the left-half plane. With proportional, lag, and lead force controllers, this result holds: even when the impedance at port A has right-half-plane poles, the impedance at port B does not.

This result indicates that the force transducer dynamics contribute to the coupled instability. If the force transducer had zero compliance, this would be equivalent to m_2 coupling directly to the environment. In this case, the robot could couple stably to rigid constraints. Because the force transducer has finite compliance, however, the system cannot stably couple to a pure constraint. The system characterized by robot port B can couple stably to a kinematic constraint but cannot couple stably to an
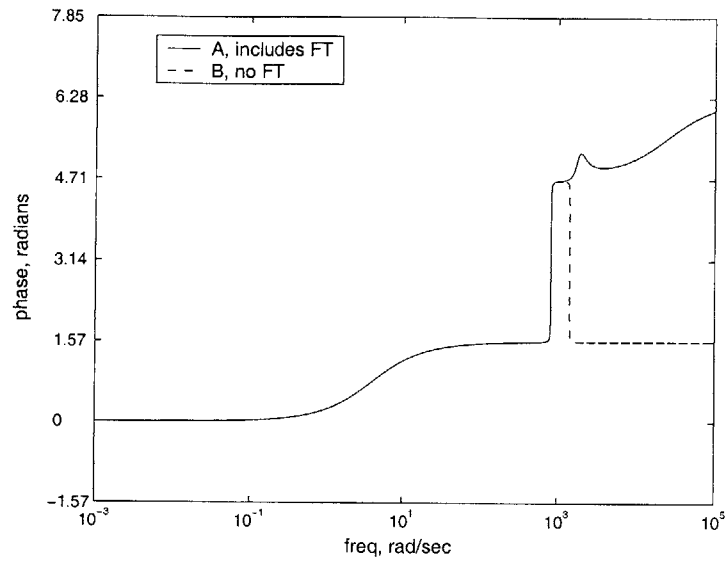
Figure A-2: Port impedance phase of single-resonance system, with and without force transducer.
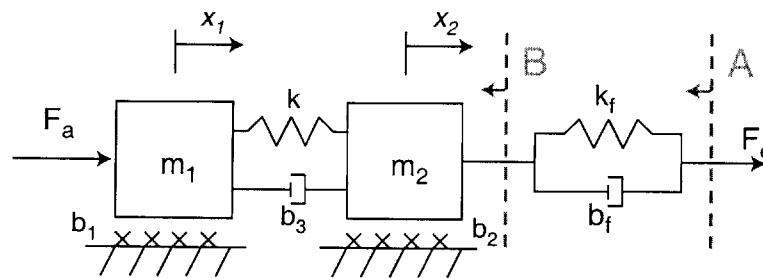

Figure A-3: System of figure 4-1. Impedance at port A includes the force transducer dynamics. Impedance at port B excludes the force transducer.

environment that consists of a stiff spring with minimal damping, as exemplified by the force transducer. This is consistent with the long-held notion that interaction with a stiff spring is the most challenging for a force-feedback system, more challenging even than interaction with a rigid constraint [27]. The difference is that the stiff spring, unlike the infinitely stiff constraint, contributes phase. For a detailed discussion of "worst-case environments" see [23].

A.3 Nonminimum phase and the small gain theorem

Finally, the presence of unstable poles in the port impedance raises some question as to the validity of applying the small gain theorem to establish the robust stability of the system shown in figure 3-6. In the example described throughout chapter 4, the system can be represented by the $M - \Delta$ structure shown in figure 4-4. To apply the small-gain theorem, either with singular values or structured singular values, it is necessary that both Δ and M be stable transfer-function matrices. In this case Δ is always stable because it is a matrix of constants. M is given by equation 4.9 and reproduced here:

$$M = \frac{\frac{1}{Z_{robot}}}{1 + \frac{m_{hn}s^2 + b_{hn}s + k_{hn}}{sZ_{robot}}} \begin{bmatrix} m_{hd}s & m_{hd} & \frac{m_{hd}}{s} \\ b_{hd}s & b_{hd} & \frac{b_{hd}}{s} \\ k_{hd}s & k_{hd} & \frac{k_{hd}}{s} \end{bmatrix} \tag{A.1}$$

Because the nominal environment (given by m_{hn}, b_{hn}, and k_{hn}) is included in the transfer function that pre-multiplies the matrix, the stability of M depends on the properties of both the port impedance and the nominal environment. The algorithm checks this fact by checking "nominal stability," the system's coupled stability when interacting with a nominal environment. Despite the fact that the robot impedance Z_{robot} has poles in the right-half plane for several of the most promising controllers, M remains stable for each of these solutions, and the computation of complementary

stability with the small-gain theorem is valid.

Although the nonminimum phase behavior of the robot port impedance is suggestive of potential stability problems, as shown here this does not predict instability in the relevant measures for this method of interaction control design.

Appendix B

Impedance shaping in fluid transmissions

The system described in chapters 6 and 7 uses a source actuator that renders target endpoint impedance effectively, and therefore the hydraulic transmission is designed to minimize its compliance, damping, and inertia. Chapter 5, however, has shown that it can be useful to place finite, carefully regulated dynamics in series between a source actuator and the environment when the actuator has high impedance. If a fluid transmission is used with such a source actuator, it may be useful to include accurately placed damping, compliance, and/or inertia in the transmission. This appendix discusses several ways to design a transmission to achieve these objectives.

B.1   Inertial impedance shapers in fluid transmissions

As shown in section 5.2.8, adding inertia can help to improve performance when force feedback is used as a control strategy to reduce apparent impedance. For certain situations, adding inertia between the force sensor and the environment can permit more aggressive force feedback that reduces the apparent impedance due to the actuator, offsetting the increased impedance due to the added mass itself. Nonzero fluid inertia

is a natural (and unavoidable) part of any transmission system that uses long, narrow transmission lines. This is shown by the model in chapter 6 as well as the test results in chapter 7. Assuming the transmission has zero compliance, this inertia produces force in proportion to the acceleration of the two pistons. Assuming that any force loop should be closed locally around the source actuator (including the transmission lines in a force loop is likely to cause problems due to transmission delays), this inertia is properly placed to serve the purpose demonstrated in chapter 5. The magnitude of this inertia can be changed by adjusting the line geometry, as discussed in chapter 6.

B.2   Series compliance and damping

In addition to inertial series dynamics, the other architecture studied in detail in chapter 5 uses a spring and damper in series between a source actuator and the environment. This structure was shown to be effective both by making the actuator more closely exhibit target behavior physically, and by permitting more aggressive force feedback on the local actuator impedance.

To create compliance in a hydraulic transmission, it is necessary to decouple the movement of the two pistons by permitting changes in the internal volume of the transmission. One way to do this is to introduce an alternative path (or *cross-flow*) for the fluid. Some mechanism must increase the internal transmission pressure as the volume increases, creating stiffness in the cross-flow. Significantly, if the damping is to be placed in parallel with the stiffness, as shown throughout chapter 5 and particularly in figure 5-4 (b_s), it must produce force in proportion to the cross-flow rate, or the *relative* velocity of the two pistons, *not* the absolute velocity of either individual piston. While the fluid line contributes damping as shown by the model in chapter 6 and the experiments in chapter 7, this is primarily damping to ground, which depends on the absolute fluid velocity, and is different from the damping shown by b_s in figure 5-4.

Although it was designed to minimize compliance, the hydraulic system prototype described in chapter 7 has finite stiffness (around 4000 N/m). This is most likely due

336

to the expansion of the flexible line with pressure, and partially to the compression of any small pockets of air that remain in the system. While this compliance might prove useful for the reasons described above, its implementation has several disadvantages. One is that it is difficult to predict, and depends on hose properties and geometry, quantities that may be dictated by other requirements on an actuator system. The other is that there is no mechanism to provide for specific parallel damping levels.

A schematic of an alternative way to introduce compliance in a fluid transmission is shown in figure B-1. Q_1 and Q_2 represent the flow associated with the movement of each of the two pistons. Q_3 represents the cross-flow. When the pistons move with equal velocity such that $Q_1 = Q_2$ and $Q_3 = 0$, there is no cross-flow, and the system operates as if the module shown is not present. If one piston is fixed such that $Q_1 = 0$, then any movement of the second piston requires a change in the internal volume of the transmission. In this design, the change occurs by the expansion or contraction of the flexible bladder. Because stress is required to deform the bladder, an increase in fluid pressure is required to expand it further. This produces stiffness in the transmission. Any flow Q_3 into or out of the bladder chamber must pass the obstructions shown. This produces losses and fluid damping. So, if $Q_1 = 0$, then the system impedance at the second piston is that of a damped spring that derives its characteristics from the compliance of the bladder and the losses from fluid flow past the obstructions. More complex operational conditions are simply combinations of these two scenarios. Any relative motion between the two pistons produces flow into the module Q_3. Any internal pressure in the transmission is provided by the expansion of the bladder. At the two pistons, the system theoretically behaves as if the two pistons are coupled by a mechanical spring and damper (in addition to any losses due to flow through the rest of the line).

B.2.1   Inertial dynamics in cross-flow

As discussed in chapter 6, fluid flow through passages results both in pressure changes related to the fluid velocity, or viscous damping, and in pressure changes related to the acceleration, or apparent fluid inertia. If the obstructions in figure B-1 produce
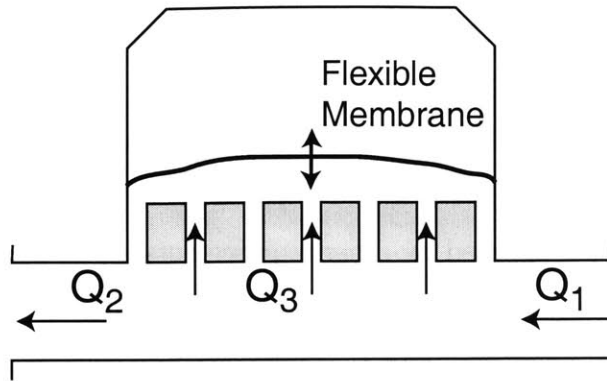
337

Figure B-1: Schematic of an arrangement to implement a hydraulic spring-damper impedance shaper. Q_1 and Q_2 represent flow at the pistons, and Q_3 is the cross-flow.

passages for the fluid to flow through, this results in apparent inertia as well as the damping discussed in the previous section. Unlike the compliance and damping introduced by the module in figure B-1, this inertia is not readily represented by an obvious mechanical equivalent (though a dynamically equivalent mechanical system can be derived). This can be demonstrated through the use of bond graphs. Figure B-2 shows schematic and bond graph representations of a mechanical system consisting of two damped masses joined by a spring and damper. While the masses m_1 and m_2 and the dampers b_1 and b_2 each produce forces related to the motion (velocity or acceleration) of one of the system ports, the series spring k_s and damper b_s produce forces related to the *relative* motion of the two ports.

Similarly, figure B-3 shows analogous representations of a fluid system consisting of the module from figure B-1 and two pipes, one on each side, that each contribute damping and inertia. The inertia and damping terms h_1, h_2 (h is used to denote fluid inertia here), b_1, and b_2 each produce pressure related to the flow at one of the system ports. The stiffness k_s, damping b_s, and inertia h_s, however, produce pressure related to the relative flow between the two system ports. The relevant difference between the bond graph in figure B-2 and that in figure B-3 is that the latter has the inertia h_s. This quantity could prove either disadvantageous or advantageous for certain applications (the search algorithm described in chapters 3-5 provides a way to test this for a particular model). In general, both quantities b_s and h_s depend on by the
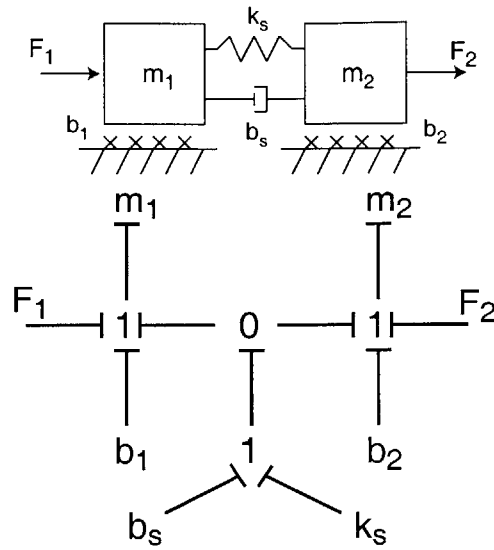
338

Figure B-2: Schematic and bond graph representations of a mechanical spring-damper between two masses.

geometry of the passages past the obstructions in figure B-1. However, it is possible to introduce damping with minimal inertia, for example by forcing fluid through small holes of negligible length. This produces a pressure drop without requiring a significant amount of fluid to accelerate faster. Several more detailed methods of introducing desired series dynamics are discussed in the next section. Because the fluid system has the flexibility to introduce inertia in this manner, it may have advantages over the simple mechanical spring-damper system.

B.2.2 Impedance shaper module concepts

In practice, it is desirable to be able to predict the dynamic characteristics of an impedance shaper while designing it. It may also be desirable for an impedance shaper to have adjustable stiffness, damping, and/or inertia. Several conceptual designs are presented for hydraulic impedance shapers that can provide variable dynamic properties. Each of the several designs permits varying degrees of independent adjustment of stiffness, damping, and inertia.
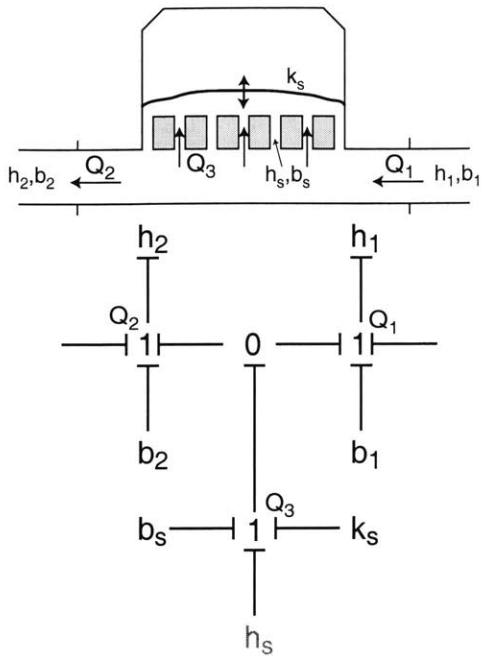
339

Figure B-3: Schematic and bond graph representations of a hydraulic impedance shaper.

Cylindrical sliding-collar impedance shaper

One way to implement series dynamics is to include a pipe segment like that shown in figure B-4 as a part of the line connecting the two pistons. The pipe has radial holes to permit cross-flow. A flexible membrane surrounds the pipe and expands when fluid under pressure fills the area outside of the pipe wall. The collars shown can be moved, as shown by the difference in their position between the two panels of figure B-4. Moving the collars from the positions shown in the first panel to the positions shown in the second panel increases both the stiffness and damping of the module. The stiffness is increased because the effective area of the flexible membrane is reduced. The damping is increased because some of the radial holes are obstructed, leaving a smaller total area for the fluid to squeeze through. If the pipe thickness is non-negligible, this also increases the apparent inertia due to flow through the radial holes. One disadvantage of this configuration is that it is difficult to change the stiffness and damping independently; one possible solution is to use
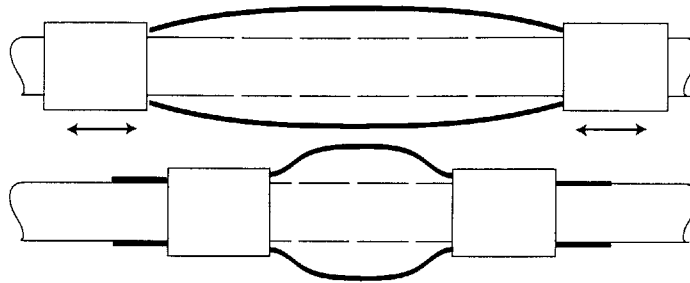
Figure B-4: A cylindrical sliding-collar impedance shaper. Moving the collars can obstruct the passages and the expansion of the bladder, changing the stiffness and damping.

an irregular pattern of holes in the pipe such that different combinations of collar locations produce different stiffness and damping values.

Variable passages to stiffness shaper

An alternative module design provides a network of passages for the cross-flow that lead to a separate "stiffness shaper" module; stiffness shapers are discussed in the next section. The network of passages, as shown in figure B-5, consists of hose segments and valves, and might consist of one or more than one pathway. If a valve is partially closed, this creates an obstruction that increases damping but does not significantly change inertia. The inertia can be varied primarily by adding and removing pipe segments with various length and diameter, or by closing off entire passages. Thus some degree of independent variation of inertia and damping can be obtained. The inertia and damping of the hoses can be determined via parallel combinations of the equations for single-pipe flow provided in chapter 6. The additional losses from the valves can be determined from experimental data (e.g. see [44]), or from direct testing.

Stiffness shaper concepts

A stiffness shaper, for use with a design like that shown in figure B-5, is a chamber that produces an increase in the contained fluid pressure as the volume of this chamber expands. Ideally for this application, the stiffness should be adjustable. Sev-
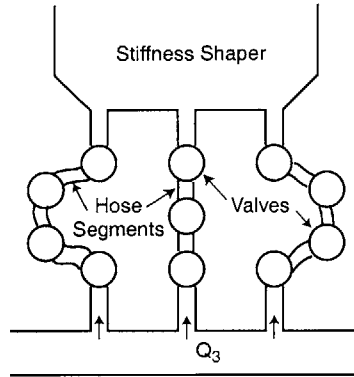
341

Figure B-5: Impedance shaper concept with variable passage geometry for adjusting damping and inertia, and separate stiffness shaper.
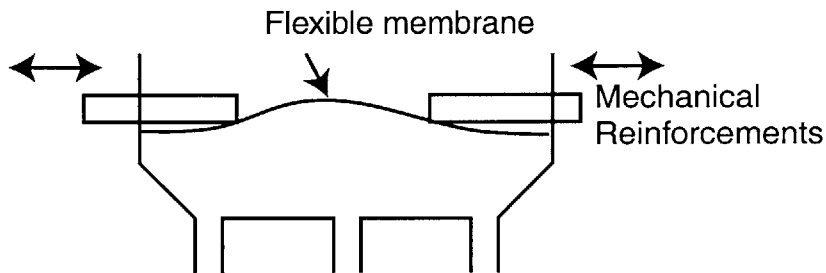
Figure B-6: Flexible membrane stiffness shaper. Movement of mechanical reinforcements changes membrane geometry and therefore stiffness.

eral different concepts are proposed here and illustrated schematically in figures B-6 through B-8.

A flexible bladder, possibly made of synthetic rubber or a similar material, provides a variable volume and increasing pressure as the volume expands, resulting from the stresses required to deform the bladder. One example of a design incorporating a bladder is shown in figure B-6. In this design, mechanical reinforcements can be positioned to prevent deformation of certain areas of the bladder, to adjust the apparent stiffness. In general, the smaller the area that is free to deform, the stiffer the system will be (this is highly dependent on geometry). The geometry pictured is only one possibility; alternative bladder geometries (e.g. cylindrical, spherical) may provide greater strength or more predictable stress-strain behavior and therefore be more easily modeled.

342

Figure B-7: Stiffness shaper concept that uses a compressible fluid to provide compliance.

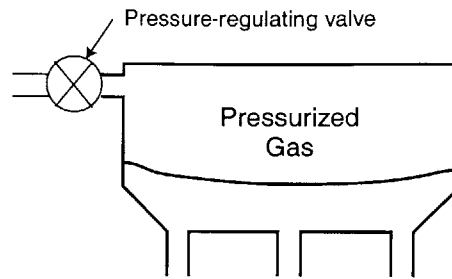Rather than using the stress in a solid bladder to generate increased fluid pressure, another possibility is to use a second fluid to provide resistance to the volumetric expansion of the hydraulic transmission. The second fluid might be a compressible gas, as shown in figure B-7. Because compressible fluids generally have a nonlinear relationship between pressure and specific volume (the ideal gas law shows this, for example), the apparent stiffness can be adjusted by altering the pressure of the gas. This design would require a separate means of regulating the pressure of the second fluid, but could permit finely tunable stiffness. A similar, more simple method is to use a simple closed tank as the stiffness shaper, and to trap a specific amount of air inside it with the working fluid. The more air that is trapped, the more compliant the system becomes. If the stiffness shaper is held at the highest elevation of the fluid system, the air, and therefore the system compliance, should stay trapped within it and avoid seeping into the fluid lines.

An alternative way to provide designed stiffness is to use a mechanical spring, as shown in figure B-8. The stiffness shaper could use a reinforced bladder or any other type of moving seal; essentially the stiffness shaper is a third piston-cylinder device that provides an opportunity to mechanically regulate the stiffness of the transmission. One or more mechanical springs can be used to adjust stiffness to the desired level.

Any of these designs, or other ideas like them, can be used to modulate the stiffness of a passive hydraulic transmission. If used in conjunction with a network like that shown in figure B-5, the stiffness, damping, and inertia can all be regulated. If made
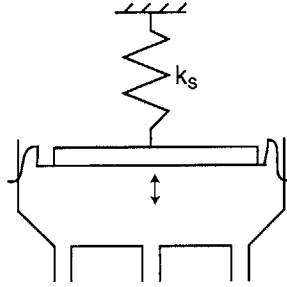
343

Figure B-8: Stiffness shaper concept that uses a mechanical spring to provide compliance.

adjustable, these can help to provide a range of possible dynamic characteristics for the transmission, so that it can fulfill the promise of impedance shapers as described in chapter 5, while also transmitting power, force, motion and impedance remotely.

B.3   Hydraulic system as integrated transmission, gear reduction, and impedance shaper

As discussed in chapters 6 and 7, a passive hydraulic system in a flexible package (such as a hose) can act as a remote transmission, transporting actuator behavior to a lightweight endpoint package. This system can also act as a geartrain by using differently-sized fluid-to-mechanical conversion elements (such as pistons) at the two ends of the system. Compared to mechanical gearing, this method is relatively free of backlash, friction, and noise.

Here it has been shown additionally that a similar hydraulic system can be used as an impedance shaper, providing advantageous dynamics between actuator and environment. Potentially a transmission can be designed that fills all three of these needs in a simple, dynamically clean (free from friction and backlash) package. Such an arrangement could permit effective high-force, low-impedance actuation in a lightweight package, using source actuators that need neither high force capacity nor low endpoint impedance. There are practical limits to this strategy, as shown throughout this thesis, but this hydraulic architecture presents the possibility of improving performance

344

and accommodating actuators that would by themselves be too heavy, too weak, or have too high intrinsic impedance to be useful for high-force haptic applications.

Bibliography

[1] 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.

[2] Airpot Corporation. *http://www.airpot.com.*

[3] Andrews, J.R. and Hogan, N. (1983) Impedance control as a framework for implementing obstacle avoidance in a manipulator. In Hardt, D. and Book, W.J., eds., *Control of Manufacturing Processes and Robotic Systems*. ASME, New York.

[4] Andriot, C. and Fournier, R. (1992) Bilateral control of teleoperators with flexible joints by the H infinity approach. In Das, H., ed., *Telemanipulator Technology, Proceedings SPIE Conference 1833*:80-91.

[5] Armstrong-Helouvry, B. (1991) *Control of Machines with Friction.* Kluwer Academic Publishers, Boston.

[6] 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.

[7] ATI Industrial Automation. *http://www.ati-ia.com.*

[8] Balas, G.J., Doyle, J.C., Glover, K., Packard, A. and Smith, R. (2001) *μ-Analysis and Synthesis Toolbox, User's Guide Version 3*. Mathworks, Natick, MA.

[9] Basdogan, C., De, S., Kim, J., Muniyandi, M., Kim, H. and Srinivasan, M.A. (2004) Haptics in minimally invasive surgical simulation and training. *IEEE Computer Graphics and Applications* **24**(2):56-64.

[10] Berkelman, P.J., Butler, Z.J. and Hollis, R.L. (1996) Design of a hemispherical magnetic levitation haptic interface device. *ASME International Mechanical Engineering Congress and Exposition* DSC **58**:483-488.

[11] Blaya, J. and Herr, H. (2004) Adaptive control of a variable-impedance ankle-foot orthosis to assist drop-foot gait. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **12**(1):24-31.

[12] The Boston Elbow. Liberty Mutual Insurance Company. *http://www.libertymutual.com/libertytechnology/be2systems.html.*

[13] Bouzit, M., Burdea, G., Popescu, G., and Boian, R. (2002) The Rutgers Master II-new design force feedback glove. *IEEE/ASME Transactions on Mechatronics* **7**(2):256-263.

[14] Brown, F.T. (2001) *Engineering System Dynamics.* Marcel Dekker, New York.

[15] Buerger, S.P. (2001) *Characterization and Control of an Interactive Robot.* SM Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[16] Buerger, S.P., Krebs, H.I. and Hogan, N. (2001) Characterization and control of a screw-driven robot for neurorehabilitation. *IEEE Conference on Control Applications/ International Symposium on Intelligent Control.*

[17] Burdet, E., Osu, R., Franklin, D.W., Milner, T.E. and Kawato, M. (2001) The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* **414**:446-449.

[18] Burgar, C.G., Lum, P.S., Shor, P.C., Van der Loos, M. (2000) Development of robots for rehabilitation therapy: The Palo Alto VA/Stanford experience. *Journal of Rehabilitation Research and Development* **37**(6):663-673.

[19] Chapel, J.D. and Su, R. (1992) Coupled stability characteristics of nearly passive robots. *Proceedings 1992 IEEE International Conference on Robotics and Automation*:1342-1347.

[20] Charnnarong, J. (1991) *The Design of an Intelligent Machine for Upper-limb Physical Therapy.* SM Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[21] Chen, E. (1999) Six degree-of-freedom haptic system for desktop virtual prototyping application. *Proceedings First International Workshop on Virtual Reality and Prototyping*:97-106.

[22] Chiaverini, S., Siciliano, B. and Villani, L. (1999) A survey of robot interaction control schemes with experimental comparison. *IEEE/ASME Transactions on Mechatronics* **4**(3):273-285.

[23] Colgate, J.E. (1988) *The Control of Dynamically Interacting Systems.* PhD Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[24] Colgate, J.E. (1994) Coordinate transformations and logical operations for minimizing conservativeness in coupled stability criteria. *ASME Journal of Dynamic Systems, Measurement, and Control* **116**:643-649.

[25] Colgate, J.E. and Brown, J.M. (1994) Factors affecting the Z-width of a haptic display. *Proceedings IEEE International Conference on Robotics and Automation* 4:3205-3210.

[26] Colgate, J.E. and Hogan, N. (1989) An analysis of contact instability in terms of passive physical equivalents. *IEEE International Conference on Robotics and Automation* **1**:404-409.

[27] Colgate, J.E. and Hogan, N. (1988) Robust control of dynamically interacting systems. *International Journal of Control* **48**(1):65-88.

[28] Colombo, G., Joerg, M., Schreier, R., Dietz, V. (2000) Treadmill training of paraplegic patients using a robotic orthosis. *Journal of Rehabilitation Research and Development* **37**(6):693-700.

[29] Copley Controls Corporation. *http://www.copleycontrols.com.*

[30] Dahleh, Mo., Dahleh, Mu. and Verghese, G. (2002) *Lectures on Dynamic Systems and Control.* Course notes for MIT course 6.241.

[31] Danaher Corporation, Kollmorgen Goldline XT product specifications. *http://www.motionvillage.com/products/motors/dc_brushless_motors/XT.html.*

[32] De Schutter, J., Bruyninckx, H., Zhu, W. and Spong, M.W. (1997) Force control: a bird's eye view. *IEEE CSS/RAS International Workshop on "Control Problems in Robotics and Automation: Future Directions."*

[33] DiaCom Corporation. *http://www.diacom.com.*

[34] Doeringer, J.A. (1999) *An Investigation into the Discrete Nature of Human Arm Movements.* PhD Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[35] Dohring, M. and Newman, W.S. (2002) Admittance enhancement in force feedback of dynamic systems. *Proceedings IEEE International Conference on Robotics and Automation* 638-643.

[36] Dohring, M. and Newman, W.S. (2003) The passivity of natural admittance control implementations. *IEEE International Conference on Robotics and Automation*:3710-3715.

[37] Drake, S. (1977) *Using Compliance in Lieu of Sensory Feedback for Automatic Assembly.* ScD Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[38] Eppinger, S.D. and Seering, W.P. (1992) Three dynamic problems in robot force control. *IEEE Transactions on Robotics and Automation* **8**(6):751-758.

[39] Eppinger, S.D. and Seering, W.P. (1987) Understanding bandwidth limitations on robot force control. *Proc. IEEE International Conference on Robotics and Automation*:904-909.

[40] Fasoli, S.D., Krebs, H.I., Stein, J., Frontera, W.R., Hughes, R. and Hogan, N. (2004) Robotic therapy for chronic motor impairments after stroke: follow-up results. *Archives of Physical Medicine and Rehabilitation* 85(7):1106-1111.

[41] Fasse, E.D., Hogan, N., Gomez, S.R. and Mehta, N.R. (1994) A novel variable mechanical-impedance electromechanical actuator. In Radcliffe, C., ed., *Dynamic Systems and Control 1994 Volume 1*. ASME, New York.

[42] Ferraro, M., Palazzolo, J.J., Krol, J., Krebs, H.I., Hogan, N. and Volpe, B.T. (2003) Robot aided sensorimotor arm training improves outcome in patients with chronic stroke. *Neurology* 61:1604-1607.

[43] Fite, K.B., Mitchell, J.E., Barth, E.J. and Goldfarb, M. (2004) Sliding mode control of a direct-injection monopropellant-powered actuator. *Proceedings American Control Conference* 5:4461-4466.

[44] Fox, R.W. and McDonald, A.T. (1973) *Introduction to Fluid Mechanics*. John Wiley & Sons, New York.

[45] Franklin, G.F., Powell, J.D. and Emami-Naeini, A. (1994) *Feedback Control of Dynamic Systems*. Addison-Wesley, New York.

[46] Gillespie, R.B., Colgate, J.E. and Peshkin, M.A. (2001) A general framework for cobot control. *IEEE Transactions on Robotics and Automation* 17(4):391-401.

[47] Gillespie, R.B. and Cutkosky, M.R. (1996) Stable user-specific haptic rendering of the virtual wall. *Proceedings ASME Dynamic Systems and Control Division* 58:397-406.

[48] Goldfarb, M., Barth, E.J., Gogola, M.A. and Wehrmeyer, J.A. (2003) Design and energetic characterization of a liquid-propellant-powered actuator for self-powered robots. *IEEE/ASME Transactions on Mechatronics* 8(2):254-262.

[49] Gomi, H. and Kawato, M. (1996) Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* **272**:117-120.

[50] Gomi, H. and Osu, R. (1998) Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. *Journal of Neuroscience* **18**(21):8965-8978.

[51] Grimes, W. (2004) Self-propelled, with a mission: clean house. *The New York Times*, July 15, 2004.

[52] Hannaford, B. and Ryu, J.H. (2002) Time-domain passivity control of haptic interfaces. *IEEE Transactions on Robotics and Automation* **18**:1-10.

[53] Hayward, V. and Astley, O.R. (1996) Performance measures for haptic interfaces. In Giralt, G., Hirzinger, G., eds., *Robotics Research: The 7th International Symposium*, Springer-Verlag:195-207.

[54] Henri, P.D., Hollerbach, J.M. and Nahvi, A. (1998) An analytical and experimental investigation of a jet pipe controlled electropneumatic actuator. *IEEE Transactions on Robotics and Automation* **14**(4):601-611.

[55] Hibbeler, R.C. (1997) *Mechanics of Materials*. Prentice-Hall, Upper Saddle River, NJ.

[56] Hogan, N. (1989) Controlling impedance at the man/machine interface. *Proceedings 1989 IEEE International Conference on Robotics and Automation*:1626-1631.

[57] Hogan, N. (1985) Impedance control: An approach to manipulation. *ASME Journal of Dynamic Systems, Measurement and Control* **107**:1-24.

[58] Hogan, N. (1990) Mechanical impedance of single- and multi-articular systems. In Winters, J. and Woo, S., eds., *Multiple Muscle Systems: Biomechanics and Movement Organization*. Springer-Verlag, New York.

[59] Hogan, N. (1988) On the stability of manipulators performing contact tasks. *IEEE Journal of Robotics and Automation* 4(6):677-686.

[60] Hogan, N. (1987) Stable execution of contact tasks using impedance control. *Proceedings IEEE International Conference on Robotics and Automation*:1047-1054.

[61] Hogan, N. and Buerger, S.P. (2005) Impedance and interaction control. In Kurfess, T., ed. *Robotics and Automation Handbook*, CRC Press, New York.

[62] Hogan, N., Krebs, H.I., Sharon, A. and Charnnarong, J. (1995) Interactive robot therapist. U.S. Patent #5,466,213.

[63] Hollerbach, J.M. (2000) Some current issues in haptics research. *Proceedings 2000 IEEE International Conference on Robotics and Automation*:757-762.

[64] Hollerbach, J.M., Hunter, I.W. and Ballantyne, J. (1992) A comparative analysis of actuator technologies for robotics. In Khatib, O., Craig, J.J. and Lozano-Perez, T., eds., *The Robotics Review 2*, MIT Press, Cambridge, MA.

[65] Hunter, I.W. and Kearney, R.E. (1982) Dynamics of human ankle stiffness: variation with mean ankle torque. *Journal of Biomechanics* 15:747-752.

[66] Jacobsen, S.C., Smith, F.M., Iversen, E.K. and Backman, D.K. (1990) High performance, high dexterity, force reflective teleoperator. *Proc. 38th Conference on Remote Systems Technology* 3:180-185.

[67] Jugenheimer, K. (2001) *A Robot for Hand Rehabilitation*. SM Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[68] Kandel, E.R., Schwartz, J.H. and Jessell, T.M. (2000) *Principles of Neural Science*. McGraw-Hill, New York.

[69] Kazerooni, H. and Guo, J. (1993) Human extenders. *ASME Journal of Dynamic Systems, Measurement and Control* 115:281-290.

[70] Kazerooni, H. (1996) The human power amplifier technology at the University of California, Berkeley. *Robotics and Autonomous Systems* **19**:179-187.

[71] Klute, G.K., Czerniecki, J. and Hannaford, B. (1998) Powering lower limb prosthetics with muscle-like actuators. *Proc. 1st National Meeting, Veterans Affairs Rehab. R & D Service.*

[72] Krebs, H.I. (1997) *Robot-Aided Neurorehabilitation and Functional Imaging.* PhD Thesis, Department of Ocean Engineering, Massachusetts Institute of Technology.

[73] Krebs, H.I., Celestino, J., Williams, D., Ferraro, M., Volpe, B.T. and Hogan, N. (2004) A wrist extension to MIT-MANUS. In Bien, Z., Stefanov, D., eds., *Advances in Human-Friendly Robotic Technologies for Movement Assistance/ Movement Restoration for People with Disabilities (Lecture Notes in Control and Information Sciences).* Springer-Verlag, New York.

[74] Krebs, H.I., Ferraro, M., Buerger, S.P., Newbery, M.J., Makiyama, A., Sandmann, M., Lynch, D., Volpe, B.T. and Hogan, N. (2004) Rehabilitation robotics: Pilot trial of a spatial extension for MIT-MANUS. *Journal of NeuroEngineering and Rehabilitation* **1**:5.

[75] Krebs, H.I., Volpe, B.T., Aisen, M.L. and Hogan, N. (2000) Increasing productivity and quality of care: Robot-aided neurorehabilitation. *Journal of Rehabilitation Research and Development* **37**(6):639-652.

[76] Krebs, H.I., Volpe, B.T., Ferraro, M., Fasoli, S., Palazzolo, J., Rohrer, B., Edelstein, L. and Hogan, N. (2002) Robot-aided neurorehabilitation: From evidence-based to science-based rehabilitation. *Topics in Stroke Rehabilitation* **8**(4):54-70.

[77] Lawrence, D.A. (1988) Impedance control stability properties in common implementations. *Proceedings 1988 IEEE International Conference on Robotics and Automation* **2**:1185-1190.

[78] Lawrence, D.A. (1993) Stability and transparency in bilateral teleoperation. *IEEE Transactions on Robotics and Automation* 9(5):624-637.

[79] Levine, W.S., ed. (1996) *The Control Handbook*. CRC Press, Boca Raton, FL.

[80] Massie, T.H. and Salisbury, K. (1994) The PHANTOM haptic interface: A device for probing virtual objects. *Proc. ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems.*

[81] Moog, Inc. *http://www.moog.com.*

[82] Moromugi, S., Izumi, K., Yoshimochi, A., Ishimatsu, T., Tanaka, T. and Feng, M.Q. (2002) Device for assisting grasping function. *Proc. 17th International Conference on Control, Automation, and Systems:*1250-1254.

[83] Morrell, J.B. and Salisbury, J.K. (1996) Performance measurements for robotic actuators. *Proceedings ASME Dynamics Systems and Control Division* 58:531-537.

[84] Nagarkatti, S.P. and Dawson, D.M. (2005) Force/impedance control for robotic manipulators. In Kurfess, T., ed., *Robotics and Automation Handbook*, CRC Press, New York.

[85] Nahvi, A., Hollerbach, J.M., Xu, Y. and Hunter, I.W. (1994) An investigation of the transmission system of a tendon driven robot hand. *Proceedings IEEE/RSJ/GI International Conference on Intelligent Robots and Systems* 1:202-208.

[86] Newbery, Miranda. (2004) Unpublished report. Massachusetts Institute of Technology.

[87] Newman, W.S. (1992) Stability and performance limits of interaction controllers. *ASME Journal of Dynamic Systems, Measurement and Control* 114(4):563-570.

[88] Newman, W.S. and Zhang, Y. (1994) Stable interaction control and coulomb friction compensation using natural admittance control. *Journal of Robotic Systems* **11**(1):3-11.

[89] NSK Corporation. *http://www.am.nsk.com*.

[90] Oak Ridge National Laboratory, Robotics and Energetic Systems Group. *http://www.ornl.gov/sci/engineering_science_technology/roboticsenergetics/exoskeletons.htm*.

[91] Palazzolo, J. (2005) *Robotic Technology to Aid and Assess Recovery and Learning in Stroke Patients*. PhD Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology. (in press)

[92] Perreault, E.J., Kirsch, R.F. and Crago, P.E. (2001) Effects of voluntary force generation on the elastic components of endpoint stiffness. *Experimental Brain Research* **141**:312-323.

[93] Perreault, E.J., Kirsch, R.F. and Crago, P.E. (2002) Voluntary control of static endpoint stiffness during force regulation tasks. *Journal of Neurophysiology* **87**:2808-2816.

[94] Peshkin, M. and Colgate, J.E. (1999) Cobots. *The Industrial Robot* **26**(5):335.

[95] Pratt, G.A. and Williamson, M. (1995) Series elastic actuators. *Proceedings, 1995 IEEE/RSJ International Conference on Human Robot Interaction and Cooperative Robots* **1**:399-406.

[96] Qian, H.P. and DeSchutter, J. (1992) Introducing active linear and nonlinear damping to enable stable high gain force control in case of stiff contact. *Proceedings 1992 IEEE International Conference on Robotics and Automation* **2**:1374-1379.

[97] Raanes, C.A. and Bodduluri, M. (2005) Robotics in medical applications. In Kurfess, T., ed., *Robotics and Automation Handbook*, CRC Press, New York.

[98] Roberts, M. (2004) *A Robot for Gait Rehabilitation.* SM Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[99] Robinson, D.W. and Pratt, G.A. (2000) Force controllable hydro-elastic actuator. *Proceedings 2000 IEEE International Conference on Robotics and Automation* 1321-1327.

[100] Robinson, D.W., Pratt, J.E., Paluska, D.J. and Pratt, G.A. (1999) Series elastic actuator development for a biomimetic walking robot. *Proceedings 1999 IEEE/ASME International Conference on Advanced Intelligent Mechatronics* 561-568.

[101] Rosen, J., Hannaford, B., MacFarlane, M.P., and Sinanan, M.N. (1999) Force controlled and teleoperated endoscopic grasper for minimally invasive surgery - experimental performance evaluation. *IEEE Transactions on Biomedical Engineering* **46**(10):1212-1221.

[102] Russo, M. and Tadros, A. (1992) Controlling dissipative magnetic particle brakes in force reflective devices. *Advances in Robotics, ASME DSC* **42**:63-70.

[103] Salcudean, S.E., Drexel, P.A., Ben-Dov, D., Taylor, A.J. and Lawrence, D.A. (1994) A six degree-of-freedom, hydraulic, one person motion simulator. *IEEE International Conference on Robotics and Automation* **2**:931-936.

[104] Sanchez-Pena, R.S. and Sznaier, M. (1998) *Robust Systems Theory and Applications.* John Wiley & Sons, New York.

[105] SensAble showcases haptic toolkits and devices enabling developers to add touch to mainstream applications. (2004) *CIMdata Newsletter* August 10, 2004. *http://www.cimdata.com/newsletter/2004/32/03/32.03.12.htm.*

[106] Sharon, A., Hogan, N. and Hardt, D.E. (1988) High bandwidth force regulation and inertia reduction using a macro/micro manipulator system. *IEEE International Conference on Robotics and Automation* **1**:126-132.

[107] Siciliano, B. and Villani, L. (1999) *Robot Force Control.* Kluwer Academic Publishers, Boston.

[108] Stramigioli, S., Fasse, E.D. and Willems, J.C. (2002) A rigorous framework for interactive robot control. *International Journal of Control* **75**(18):1486-1503.

[109] Stramigioli, S. and Scherpen, J.M.A. (2000) Intrinsically passive handling and grasping. *European Experts Colloquium on Packaging & Handling Automation*:14-17.

[110] Tilley, S.W. and Cannon, R.H. (1986) End point position and force control of a flexible manipulator with a quick wrist. *Proceedings AIAA Guidance, Navigation, and Control Conference*:41-49.

[111] Townsend, W.T. and Guertin, J.A. (1999) Teleoperator slave - WAM design methodology. *The Industrial Robot* **26**(3):167.

[112] Trumper, D.L. (1999) ASPE tutorial: Precision mechatronics. Unpublished.

[113] Verdirame, J. (2000) *Characterization of a Hydraulic Actuator for a Functional Magnetic Resonance Imaging Robot.* SB Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[114] Volpe, B.T., Krebs, H.I., Hogan, N., Edelsteinn, L., Diels, C.M. and Aisen, M.L. (1999) Robot training enhanced motor outcome in patients with stroke maintained over 3 years. *Neurology* **53**:1874-1876.

[115] Volpe, R. and Khosla, P. (1993) A theoretical and experimental investigation of explicit force control strategies for manipulators. *IEEE Transactions on Automatic Control* **38**(11):1634-1650.

[116] Wen, J.T. and Murphy, S. (1991) Stability analysis of position and force control for robot arms. *IEEE Transactions on Automatic Control* **36**(3):365-370.

[117] Wheeler, J. (2004) *An Ankle Robot for a Modular Gait Rehabilitation System.* SM Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.

[118] Whitney, D. (1977) Force feedback control of manipulator fine motions. *ASME Journal of Dynamic Systems, Measurement and Control* **99**:91-97.

[119] Whitney, D. (1987) Historical perspective and state of the art in robot force control. *International Journal of Robotics Research* **6**(1):3-14.

[120] Wilfinger, L.S., Wen, J.T. and Murphy, S.H. (1994) Integral force control with robustness enhancement. *IEEE Control Systems Magazine* **14**(1):31-40.

[121] Willems, J.C. (1972) Dissipative dynamical systems part I: General theory. *Archives for Rational Mechanics and Analysis* **45**(5):321-351.

[122] Williams, D.J., Krebs, H.I. and Hogan, N. (2001) A robot for wrist rehabilitation. *Proceedings 23rd Annual EMBS International Conference*:1336-1339.

[123] Wyatt, J.L., Chua, L.O., Gannett, J.W., Goknar, I.C. and Green, D.N. (1981) Energy concepts in the state-space theory of nonlinear n-ports: Part I - passivity. *IEEE Transactions on Circuits and Systems* **28**(1):48-61.

[124] Zinn, M., Roth, B., Khatib, O. and Salisbury, J.K. (2004) A new actuation approach for human friendly robot design. *International Journal of Robotics Research* **23**(4-5):379-398.