CHARACTERIZING AND CONTROLLING THE HIGH-FREQUENCY DYNAMICS OF HAPTIC INTERFACES

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF
MECHANICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Katherine Julianne Kuchenbecker

June 2006

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Günter Niemeyer Principal Advisor

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Mark R. Cutkosky

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

J. Kenneth Salisbury, Jr.

Approved for the University Committee on Graduate Studies.

Abstract

Humans are amazingly adept at eliciting and interpreting touch-based feedback during interactions with everyday objects, naturally leveraging this wealth of dynamic information to guide both exploratory and dexterous manipulation. Haptic interfaces attempt to recreate the feel of real objects for telerobotic or virtual interactions, allowing the user to touch distant or unreachable environments and computer-generated models through a lightweight robotic arm. Current haptic rendering techniques, which use position or force feedback, generally cannot convey a crisp contact with a hard object, nor can they convey the fine features of a textured surface; instead, portrayed objects feel overly soft and unnaturally smooth or oscillatory, limiting the usefulness of teleoperation systems and virtual environments.

This work has developed methods for stably endowing impedance-type haptic interfaces with the high-frequency (20 to 1000 hertz) feedback signals necessary to make virtual and remote objects feel nearly indistinguishable from their real counterparts. The fundamental insight for this undertaking is that haptic systems have internal electrical, mechanical, and biomechanical dynamics that strongly influence their performance. Indeed, simply replaying a scaled version of a recorded acceleration as a motor current command does not produce the intended acceleration at the user's hand. These dynamics can be modeled by careful application of developed identification techniques, including comprehensive evaluation and successive isolation, and the resulting models can be used to improve interaction realism in two main ways.

First, the controller can precisely create high-frequency fingertip accelerations during contact with remote or virtual objects by inverting the interface's dynamics before playback. For teleoperation the target accelerations are measured in real time at the remote manipulator, and in virtual environments they are pre-recorded. When overlaid with position-based feedback, high-frequency acceleration matching creates vibrations at the user's hand that

closely correspond to the specified signals and that feel almost identical to the real object, as confirmed by a human subject study.

Second, the dynamic relationship between haptic feedback command and measured device position can be estimated and canceled to improve the stability of a haptic interface. A position-force teleoperation system that vibrates unnaturally during contact with hard objects behaves well if induced master motion is canceled from the remote robot's movement command. Canceling induced master motion stabilizes strong high-frequency feedback and allows the user to feel the remote environment more clearly.

Both of these control strategies use a dynamic model of the haptic interface's high-frequency behavior to make remote and virtual interactions feel more real. Application of these techniques to minimally invasive surgery and medical simulation is specifically promising, as it would allow physicians to feel the hardness and texture of the structures being manipulated, potentially facilitating new procedures and improving patient outcomes.

Acknowledgments

Many individuals and organizations have supported me in the creation of this dissertation, contributing invaluably to my research, writing, and state of mind. My advisor, Professor Günter Niemeyer, has played a pivotal role on all three of these fronts ever since our first meeting in January of 2002. His quick mind, excellent technical skills, and infectious enthusiasm inspired me to pursue my Ph.D. at Stanford, and our countless interactions over the intervening years have helped me know that I chose a wonderful path. Someday I hope to be able to train my own graduate students as expertly and gently as Günter has trained me.

My evolution as a doctoral student has also benefited greatly from the influence of Professor Mark Cutkosky, as he co-advised the first major research project in which I participated. Although it does not appear in this thesis, the contact display work that I did with Mark, Günter, and Will Provancher helped shape me as a researcher, introducing me to the field of psychophysics and to the art of technical writing. Mark has aided my later research by graciously loaning a pair of force-feedback joysticks, and the body of work he initiated with Rob Howe and Allison Okamura has served as a source of inspiration throughout my own investigations. Professor Ken Salisbury has played a similarly important role, giving and lending many useful pieces of equipment to our lab and keeping a watchful eye over our progress. I sincerely appreciate the input that Günter, Mark, and Ken have provided on this thesis; their comments have strengthened and clarified the research and writing throughout.

I never would have met Günter or decided to stay for a doctorate if not for the guidance of Professor Sheri Sheppard. On an auspicious day in January 2002, she patiently taught me about the inner workings of graduate school, listened to my interests, and suggested that I talk with the department's newest faculty member, an expert in haptics and teleoperation. I joined Günter's Telerobotics Lab simultaneously with June Park and Neal Tanner: along

with Jonathan Fiene and Probal Mitra, we have formed the core of our advisor's new research group, building from a few computers in the basement of Ginzton to become a vibrant, well-established laboratory. I have tremendously enjoyed the friendship of these four fellow students, and I wish them all the best as they finish their own degrees and head out into the world. I am honored to have worked alongside them and the many other students who have helped create the Telerobotics Lab over the last four and a half years.

The technical inspiration for this thesis stemmed from the excellent system identification course of Professor Chris Gerdes; the final class project that June Park, Valerie Seymour, and I conducted yielded our lab's first conference paper and planted a seed in my mind about the importance of the internal dynamics of haptic interfaces. My graduate studies were generously supported by a research fellowship from the National Science Foundation and a scholarship from the Achievement Rewards for College Scientists Foundation. Research funding from the Stanford Office of Technology Licensing, Intuitive Surgical, Inc., and the National Science Foundation also instrumentally supported my work, as did several Stanford teaching assistantships. I especially enjoyed my two years as a teaching assistant in the Product Realization Lab; the design, manufacturing, and teaching experience afforded by that time has been and will continue to be invaluable to my identity as an engineer.

The last group of individuals to whom I am deeply grateful has supported my personal happiness and state of mind throughout graduate school. I would have been terribly lonely without the Mechanical Engineering Women's Group, and I appreciate the wisdom imparted by our many Thursday afternoon speakers. Bree Sharratt has been my kindred soul, and I shall always treasure her perspectives and her company over a cup of tea. Outside the lab, Jonathan has been the perfect partner; I would never have reached this point without his unflagging support. Nor would I be who I am today without the guidance and love of my family, especially my parents, Drs. Stephen and Shari Kuchenbecker. I am proud to become the third Dr. Kuchenbecker in the family, and I look forward to supporting my brother and sister in their own quests for this unique distinction.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The primary manner in which humans can affect their surroundings is through physical contact. We use our muscles and bones to move around the world, exerting forces and torques on various objects to accomplish tasks like brushing our teeth, driving a car, and assembling a new piece of furniture. We would be nearly helpless without this ability to change the physical organization of the world around us, an ability that is strongly mediated by our capacity to sense the effects of manual interactions. Indeed, almost all living organisms have a sense of touch, the capacity to detect contact between self and environment [6]. It is the oldest and most primitive of senses, but its intricacies are far from understood [42].

When was the last time you found yourself without a sense of touch, and how did this loss affect you? Perhaps your hands became numb from cold and you found it hard to unfasten your jacket. Maybe your leg fell asleep while you were sitting in a peculiar position, and you experienced difficulty walking afterwards. Or perhaps you received a shot of local anesthetic from the dentist and struggled to talk and eat until its numbing effect had diminished. Each of these situations put you in the uncomfortable position of being unable to detect the interactions between part of your body and its physical environment. Without such feedback, you are forced to rely more heavily on visual and auditory input, if it is available, and simple tasks become much more laborious and challenging. The sense of touch plays a pivotal role in the human ability to interact effectively with the physical world, but its natural ubiquity often causes us to take it for granted [11].

The sense of touch also plays a central role in many professional activities. If you were undergoing heart surgery, you would want the surgeon to have full sensation in his or her

fingers. If you were paying a mechanic by the hour to work on your car, you would want him or her to be adept at disassembling and reassembling the components of your engine by hand. And if a technician was working to defuse a bomb hidden beneath a bridge near your house, you would want him or her to be able to distinguish between different wires and components by feel alone. Experts like these undergo significant training to learn to interpret touch-based feedback and accomplish a wide variety of important, difficult tasks for society. For further understanding of touch-based interactions, a detailed treatment of human sensory and motor capabilities will be provided in Section 1.1.

The adjective "haptic" is a more formal synonym for the term "touch-based." It stems from the Greek word *haptikos*, which means "able to touch or grasp," a descendant of the Greek verb *haptein* or "fasten" [34]. Haptic feedback is often taken to include both tactile sensation from the skin and force sensation from the muscles. As technology has improved, humans have developed systems that enable haptic interaction with distant or virtual environments, rather than just real objects. As described in Section 1.2, these so-called haptic interfaces attempt to stimulate the user's sense of touch in a manner that resembles real interactions, allowing users to conduct activities from afar or to practice manual interactions through simulation. Teleoperation and virtual environments are discussed in Sections 1.3 and 1.4, respectively, and their existing performance limitations are highlighted in Section 1.5.

This thesis develops methods for making haptic interfaces to remote and virtual interactions feel more realistic, as presented in Section 1.6. When there is close correspondence between the touch-based feedback provided during teleoperation and that received during a real manipulation, the user will be able to leverage skills from real experiences and more easily perform tasks in the remote environment. Similarly, when the haptic feedback from a model-based simulation is closely matched to that of a real interaction, the user will be able to interact more naturally and learn tasks that translate well to the real world. On the other hand, when haptic feedback is a poor approximation of reality, the operator loses his or her ability to feel the effects of an interaction, and tasks become more difficult and time-consuming. To fully realize the potential of remote and virtual touch-based interactions, we must design haptic interfaces to match human capabilities and needs, as overviewed in Section 1.7 and put forth in the body of this thesis.

1.1 Human Sensory and Motor Capabilities

The human body is an amazingly sensitive and able organic machine. It is generally considered to have five main sensory input streams: vision, hearing, taste, smell, and touch, a classification that stems from Aristotle [135]. Each sense enables a unique set of measurements about the physical world, the aggregation of which guide human behavior. Unlike the four other senses, which are afforded by sensory organs located on the head alone, touch is distributed throughout the body. Furthermore, touch is made up of several different types of sensations, including pressure, vibration, temperature, shear force, and pain, each of which can be localized to a certain area of the body with varying degrees of acuity.

While vision provides us with spatial information about the world, we use the sense of touch primarily to ascertain material properties such as hardness and texture [75]. During typical manual interactions, the human hand experiences a broad spectrum of forces ranging from steady state to over one kilohertz. These signals are detected by a rich array of mechanoreceptors in the skin, muscles, and joints, naturally guiding both dexterous and exploratory interactions. There are four main types of mechanoreceptors in glabrous (non-hairy) human skin [68], and their functions are summarized in Table 1.1.

As delineated by Johansson in 1976, the mechanoreceptors in the glabrous skin of the human hand can be defined by the speed of their adaptation to mechanical stimuli and the size of their receptive field. Fast-adapting (FA) mechanoreceptors react primarily to changes in stimulus state, and they do not exhibit a prolonged response to a steady input; in contrast, slow-adapting (SA) mechanoreceptors respond to static stimuli like a constant deflection of the skin [71]. Mechanoreceptors are further differentiated by whether they

Table 1.1: Mechanoreceptors in the glabrous skin of the human hand.

		Receptive Field Size	
		Small	Large
Adaptation Speed	Fast	FA I Meissner Corpuscle 8 to 64 Hz Spatial Deformation	FA II Pacinian Corpuscle Greater than 64 Hz Non-localized Vibration
	Slow	SA I Merkel Complex 2 to 32 Hz Spatial Curvature	SA II Ruffini Ending DC to 8 Hz Directional Stretch

have a small receptive field with well defined borders (Type I) or a large receptive field with diffuse borders (Type II) [71]. Mechanoreceptors with a small field are generally located close to the surface of the skin, and those with a large field occur more deeply.

The four types of mechanoreceptors have been documented extensively via microneurography, and they have been associated with the four end organs (sensor units) listed in Table 1.1. The frequency range in which each mechanoreceptor type is most sensitive is also listed, taken from [70], though it should be noted that these sensitivity bands depend somewhat on the magnitude of the input. Furthermore, the cited study tested discrete frequencies and did not establish an upper limit for the response of Pacinian corpuscles (PCs), noting that they still responded well at 400 hertz [70]. Bell et al. state the bandwidth of Pacinian corpuscles as 20 to 1000 hertz, based on an extensive review of pertinent literature [10]. They further note that these receptors respond most strongly to excitation between 250 and 550 hertz, and the peak response frequency depends on the sensitivity criterion used, commonly lying between 250 and 450 hertz.

Following their distinct response characteristics, each mechanoreceptor has been found to respond most intensely to a certain type of stimulus. Most interestingly, the Pacinian corpuscles (FA II) are known to respond best to the high-frequency vibrations that stem from contact between hard objects. Westling and Johansson conducted a relevant study in which subjects repeatedly lifted and set down a small test object while the responses from single tactile units were recorded: these researchers found that the "FA II units responded distinctly to the mechanical transients associated with the start of the vertical movement and especially with the sudden cessation of movement at the terminal table contact" [154]. As Pacinian corpuscles are located deep within the tissue of the hand, they have large receptive fields and often respond to stimuli that are applied many centimeters away; Pacinian corpuscles are especially implicated in tool- and probe-mediated interactions [12], providing humans with ample information about macroscopic material properties as well as fine surface features.

In contrast to our high-bandwidth sensory system, we humans cannot move or position our fingertips at frequencies above eight to ten hertz [92, 136]. This asymmetry in sensory and motor capabilities causes humans to use low-frequency hand motions to elicit information-laden high-frequency responses from the environment; Johansson calls this strategy "sensory discrete-event driven control" and notes that it is commonly used over short time scales during manipulation of physical objects [69]. For example, consider

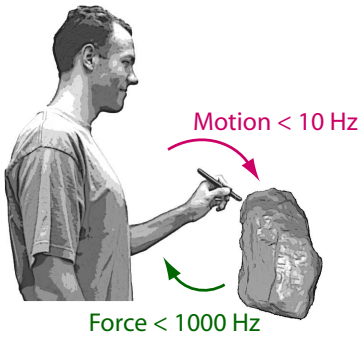

Figure 1.1: Asymmetry between human actuation and sensation.

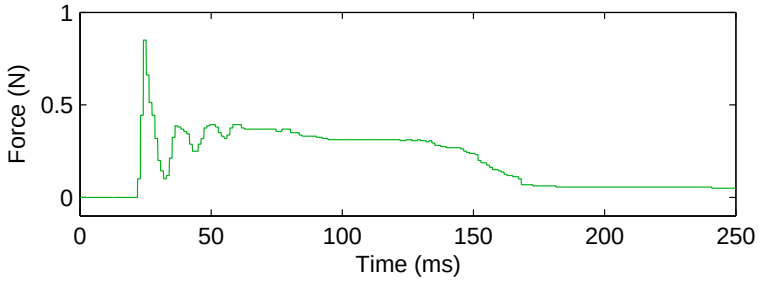

Figure 1.2: Contact force from tapping on a hard wooden table with a plastic stylus.

the act of tapping on a wooden table or a stone block with a pen, as shown in Figure 1.1. Though the person initiates the motion, it is the physical properties of the object and tool that shape the transient details of contact. Each tapping event creates sudden accelerations with frequency content up to several hundred hertz, strongly stimulating the Pacinian corpuscles in the hand and fingertips. These tool-mediated contact vibrations provide the person with far richer information about the object's material properties than could be obtained by merely pressing the stylus into it, as demonstrated by LaMotte's investigation of human capabilities during interactions with real test samples of varying softness [86].

Another way to see the role of high-frequency haptic feedback is to examine the force transient produced by contact with a hard object. Figure 1.2 shows a force signal recorded from tapping on a hard wooden surface with a stylus. This contact is seen to produce two distinct, superimposed forces: an initial high-frequency transient and a slower extended response. Over long durations, the object opposes penetration, yielding a quasi-static, low-frequency reaction force. The shape of the short-duration transient at impact is determined

by material properties and initial user conditions, including grasp configuration and incoming velocity. Impact transients generally take the form of decaying sinusoids [152], though multiple resonant modes and intermittent contact may lead to a more complex response. It is these signals, lasting tens of milliseconds, that create high-frequency accelerations and allow the user to infer the material properties of the object. Together with low-frequency forces sensed in the muscles and tendons, high-frequency accelerations provide important haptic cues that allow humans to interact easily with their physical surroundings. Haptic interfaces must provide these same high-frequency vibrations during remote and virtual interactions if they are ever to succeed at replicating the feel of real objects.

1.2 Haptic Interfaces

A haptic interface is a system that allows a human to feel as though he or she is touching a distant or virtual environment. These interfaces allow the user to move part of his or her body, typically the hand, and feel touch-based feedback such as forces, vibrations, or changing contact location. Haptic interfaces generally resemble robotic systems and include electromechanical sensors and actuators in a variety of configurations. The portrayed relationship between motion and haptic feedback is usually meant to resemble an interaction with a real object so that the user can perform the task or explore the environment in a natural manner [133]. To augment the illusion, such interfaces are almost always accompanied by visual feedback and can include auditory feedback as well.

There are two main categories of haptic interfaces: admittance type and impedance type, as described below in Sections 1.2.1 and 1.2.2. These two types differ in how they control the relationship between movement and force at the user's hand, as these two quantities are dynamically coupled. Most devices fall into one of these dual categories, though a third class that allows the user to move freely and responds with non-force-based feedback, such as low-amplitude vibrations, spatial patterns, or temperatures, exists as well. This body of research is concerned with impedance-type haptic interfaces that apply forces to the user via actuators such as DC motors. A related class of devices uses brakes rather than actuators to impede the user's motion during contact with remote or virtual objects; selecting a braking technology [106] and coordinating multiple degrees of braking [77] are both present challenges for this display type, making passive devices far less common than active displays.

1.2.1 Admittance Type

Haptic interfaces of the admittance type measure the force that the user exerts on the handle and vary the amount of motion that results. Typically constructed as strong, non-backdrivable mechanisms with a force sensor at the endpoint and highly geared motors with position sensors at the base, admittance-type haptic interfaces attempt to present the user with the following dynamic relationship,

$$\vec{v}_{hi} = Y_{env}\, \vec{F}_{user}, \tag{1.1}$$

where $\vec{v}_{hi}$ is the output velocity of the haptic interface, Y_{env} is the admittance of the remote or virtual environment, and $\vec{F}_{user}$ is the input force exerted by the user on the interface. The admittance that the user perceives may deviate from the target Y_{env} due to imperfect force sensing, imperfect velocity output, and the dynamics between the handle and the motors, which often stem from compliant members, friction, and other nonideal effects. A related class of haptic interfaces for virtual environments tracks the movement of the user's hand without contact, using optical sensors, and positions itself at the location of a virtual object's surface when contact is imminent [157]; such systems allow unrestricted free-space motion and provide firm contacts, but they have difficulty tracking fast hand movements and are thus not yet widespread.

Recent examples of admittance-type haptic interfaces are the FCS Robotics HapticMaster [31] and the Johns Hopkins University Steady-Hand Robot [143]. Systems such as these excel at inhibiting movement ($Y_{env} \approx 0$), as they can use the mechanism's natural stiffness to resist near-arbitrary user forces. In contrast, such interfaces generally struggle to let the user move freely ($Y_{env} \gg 0$), as the controller must mask the inherent mass, friction, and stiffness of the device using limited-resolution force sensing and limited-authority actuators. Although admittance-type haptic interfaces are enjoying a resurgence through the improved mechanical and control design of systems like the Northwestern University Cobotic Hand Controller [36, 37], several other features have limited their widespread use, including the high cost of good force sensors, the danger of using non-backdrivable robots near people, and the often encumbered feel of free-space motion.

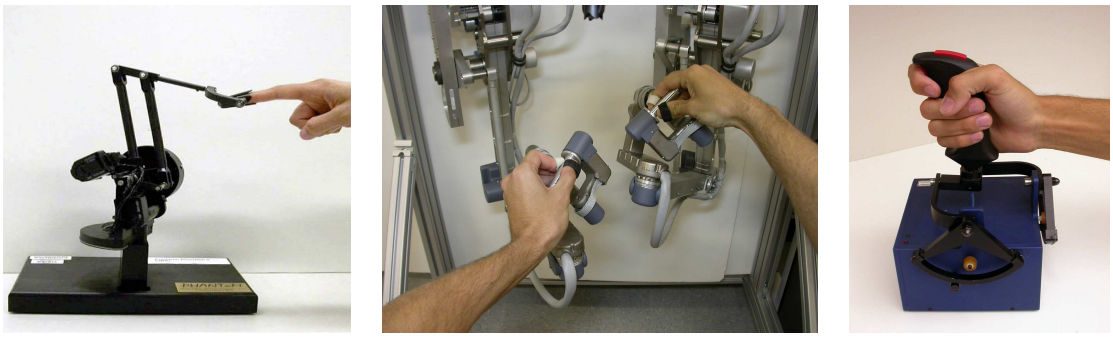

SensAble Phantom | Intuitive Masters | Immersion Joystick

Figure 1.3: Common impedance-type haptic interfaces.

1.2.2 Impedance Type

Haptic interfaces of the impedance type measure the motion that the user imposes on the handle and vary the amount of force applied in response. Typically constructed as lightweight, backdrivable mechanisms with DC motors and position sensors at the base, impedance-type haptic interfaces typically attempt to present the user with the following dynamic relationship,

$$\vec{F}_{hi} = Z_{env}\ \vec{v}_{user}, \tag{1.2}$$

where $\vec{F}_{hi}$ is the output force of the haptic interface, Z_{env} is the impedance of the remote or virtual environment, and $\vec{v}_{user}$ is the input velocity applied by the user to the interface. The impedance that the user perceives may deviate from the target Z_{env} due to imperfect velocity sensing, imperfect force output, and the dynamics between the handle and the motor, which often stem from compliant members, friction, and other nonideal effects.

The most common example of an impedance-type haptic interface is the SensAble Phantom [98], and other examples include the master manipulators of the Intuitive da Vinci system [52, 108] and the Immersion Impulse Engine 2000 [66], all of which are pictured in Figure 1.3. Systems such as these can easily allow near free-space motion ($Z_{env} \approx 0$), as the mechanism's natural dynamics possess low inertia and low friction. In contrast, these systems encounter difficulties in displaying very stiff, dissipative, or massive environments ($Z_{env} \gg 0$) as the controller must try to recreate these dynamics with limited-resolution position sensing and limited authority actuators. Impedance-type haptic interfaces are generally more widespread than admittance-type systems due to their relatively simple design, higher safety around humans, and the good haptic feel of free-space motion. The remainder

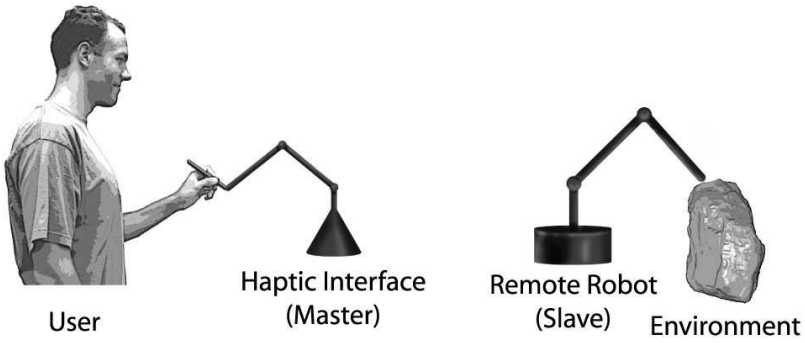

Figure 1.4: Teleoperation paradigm.

of this thesis will deal exclusively with active impedance-type devices, developing control methodologies that improve the feel of hard contact in teleoperation and virtual environments without sacrificing the natural benefits of a lightweight haptic interface.

1.3 Teleoperation

For over fifty years, teleoperation has promised users the ability to manipulate and perceive a remote environment as though it were directly accessible. As illustrated in Figure 1.4, teleoperation systems employ a robotic mechanism as the user's proxy at the remote site and use a haptic interface to measure the operator's motions and provide appropriate haptic feedback. The telerobotic system acts as an extended tool, leveraging the operator's skills and decision-making abilities into a setting beyond normal reach. Such technology enables humans to safely handle toxic waste, assemble space equipment from Earth, and explore the deep sea, all tasks that would be difficult to perform directly or to automate. Other potential applications include treating wounded soldiers on a battlefield, hand-manipulating atoms at the nanoscale, human-guided manufacturing, and assistive technology for the disabled.

Although it was originally developed for industrial applications such as nuclear decommissioning and underwater construction, telerobotic technology is seeing a strong recent interest driven by medical applications, starting with minimally invasive surgery. Here, the use of tiny robotic tools allows surgeons to operate on internal organs through ports that are just one centimeter in diameter, avoiding the large incisions of traditional surgery along with the associated patient trauma and recovery time. The surgeon guides the movements of the robotic instruments via natural hand motions while seated comfortably at a console

next to the patient [52, 129]. Ideally, the system should provide the doctor with realistic haptic feedback throughout the operation, enabling skillful tissue differentiation, incisions, and suturing.

Another modern instance of teleoperation arises when a machine's mechanical control linkage is replaced by a computer-controlled electromechanical connection, as in automotive steer-by-wire or manual operation of a computer numerical control (CNC) machining center. In both of these situations, the human needs to be able to tell the machine how to move and would also benefit from feeling forces or vibrations as they occur during the task. The computer connection is more flexible than the original mechanical linkage, but providing appropriate teleoperative feedback often proves to a challenging control-design task.

1.4 Virtual Environments

Haptic simulations aim to recreate the feel of real manipulations for virtual reality. As illustrated in Figure 1.5, a haptic interface is used to measure the user's motion, which is mapped into a virtual (simulated) world. Appropriate reaction forces are computed in real time from the model and displayed to the user using the actuators on the device. Virtual environment systems allow the user to interact with geometric and/or dynamic data that they otherwise would not be able to touch. For instance, a user can explore a complex computer-aided design (CAD) model of an automotive dashboard and assess its feel without need for a physical prototype. Researchers can feel topographical data from mountain ranges or micro-electro mechanical systems (MEMS) to gain a stronger intuition for their behavior, as an engineering student can interact with a model of a dynamic

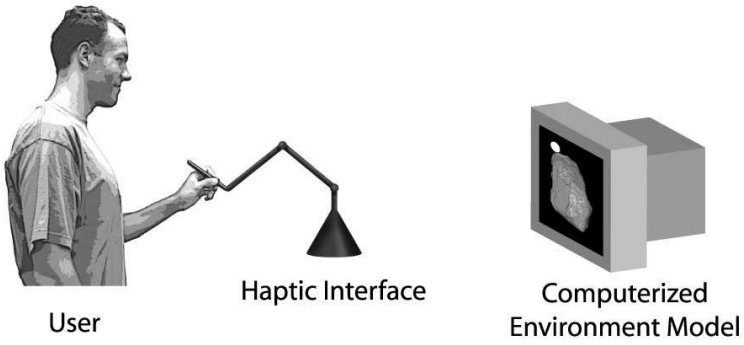

Figure 1.5: Virtual environment paradigm.

system to better understand the roles of its various parameters [116]. Other applications of virtual environments include the creation and manipulation of three-dimensional shapes for sculpture or design and the enhancement of entertainment media.

Another promising use of virtual environments is medical simulation and training [2, 26, 85, 145]. Instead of learning to perform a procedure solely by watching an expert doctor conduct it on a patient, medical students and residents can use special haptic interfaces to practice a wide variety of important treatments. The computerized models can be used to teach both cognitive and manual skills, and the system can track each individual's performance via well-defined metrics such as completion time, errors, and total applied force. Ideally, interactions such as these would be as vivid as performing the procedure on a real patient, allowing the trainee to feel the full spectrum of haptic feedback available during real operations without risking the health of a real patient or requiring the cost of an animal model or cadaver.

1.5 Hard Contact Performance Limitations

While they have been used in a wide variety of applications for both teleoperation and virtual environments, haptic interfaces have yet to achieve the levels of haptic realism required for widespread adoption [16, 25, 117]. As noted in Section 1.2.2, the primary area in which impedance-type interfaces struggle is the rendering of hard contact. Tapping a metal tool against a bone or dragging a stylus across a stone surface creates high-frequency vibrations at your fingertips that help you locate the surface and discern its material and textural properties. The vast majority of teleoperation and virtual-environment systems do not convey these high-frequency signals because of a historical reliance on closed-loop position and velocity feedback. This common control strategy effectively places a spring and damper between the motors of the haptic interface and either the motors of the slave robot or the surface of the virtual object. In teleoperation this algorithm is often called position-position control, and it can also be termed proportional-derivative (PD) feedback in both remote and virtual applications. The stiffness of the controller spring is severely limited by closed-loop stability requirements [1, 32, 33, 88], which low-pass filters contact transients and prevents high-frequency vibrations from reaching the hand of the user.

Providing only low-frequency haptic feedback requires the operator to rely unnaturally

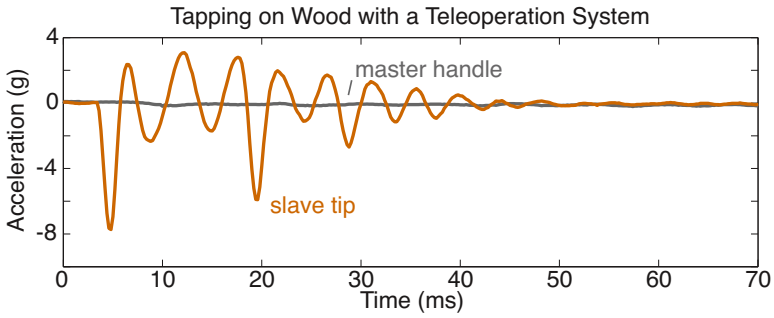

Figure 1.6: Slave tip and master handle accelerations for a telerobotic tap on a piece of wood during position-position control.

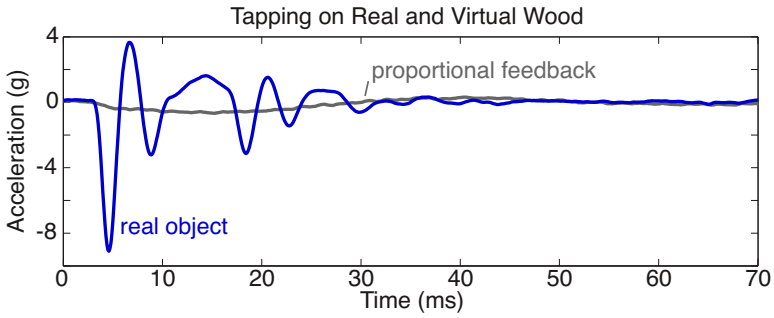

Figure 1.7: Stylus accelerations for a human tapping on a real piece of wood and on a virtual surface rendered with proportional feedback.

on vision and hearing, diminishing their sensitivity to environmental properties and prolonging tasks. When conducting surgery telerobotically, the lack of high-frequency haptic feedback makes all items feel like soft foam, and surgeons cannot adequately detect tissue reactions, suture tension, or even tool collisions [17, 20, 35, 125, 140]. Position-based control also fails to deliver haptic feedback that feels fully authentic in virtual environments, once again leaving all surfaces feeling soft and undefined [67, 117, 153]. Ideally, remote and virtual haptic interactions would be indistinguishable from real manipulations, and the lack of high-frequency vibrations is the most salient omission from present systems.

The feel of a hard contact can be quantified by examining the high-frequency accelerations it creates, since the Pacinian corpuscles in glabrous human skin serve approximately as biomechanical accelerometers [10, 12]. As an example of the poor performance of position-position feedback, Figure 1.6 plots the accelerations experienced by the remote robot's end effector and the user's hand when teleoperatively tapping on a piece of wood. While the

slave tip undergoes a high-magnitude, high-frequency acceleration transient at contact, the user feels approximately nothing. Figure 1.7 plots the corresponding comparison for virtual environments: the acceleration that stems from tapping on a real piece of wood and that created by tapping on a virtual surface with the same contact velocity. The acceleration produced by the proportional feedback is significantly lower in magnitude and frequency than that created by a tap on the real object. The poor correspondence between the two curves in each of these diagrams must be addressed and is a central focus of this dissertation.

Position-force control is another common algorithm used in teleoperation [28,56,88]. The remote robot tracks the motion of the haptic interface via PD feedback, and environmental interaction forces are explicitly measured with a force sensor at the slave's end effector. These forces are scaled up or down and displayed to the user via the actuators on the master device. While this scheme does capture the high-frequency effects of environmental contact, it fails to portray them to the user's hand for two main reasons. First, high-frequency feedback forces are distorted by the haptic interface's internal dynamics, for the connection from master motor to hand often contains several structural resonances [1].. Thus the force that the user's hand feels is not merely a scaled version of that experienced by the slave tip. Second, such feedback forces compromise closed-loop stability and generally limit the force-feedback scale factor to inadequately low values [88]. When using a teleoperator under position-force control, the user typically feels distorted, low-magnitude forces that do not adequately convey the slave's contact experience. The stability limitation encountered for high-magnitude, high-frequency haptic feedback requires consideration and is another central focus of this dissertation.

1.6 Improving Haptic Interface Performance

Haptic interfaces presently fail to portray the high-frequency hand accelerations that define hard contact, leaving all remote and virtual objects feeling soft and indistinct. The source of this problem is twofold: closed-loop position feedback is inherently limited to low frequencies, and high-magnitude, high-frequency haptic feedback like that from a slave force sensor causes closed-loop contact instability. The fundamental insight of this thesis is that both of these problems can be confronted via open-loop control adjustments that are based on a dynamic model of the haptic interface as it is held by a user.

1.6.1 Creating High-Frequency Fingertip Accelerations

Generating realistic high-frequency hand accelerations with a position-position controlled teleoperation system is nearly impossible, as the series dynamics of the slave mechanism, controller, and master mechanism low-pass filter this feedback significantly. Even a perfectly rigid mechanical connection from slave motor to master motor would not fully suffice, as the intervening robot dynamics would still distort the feedback signal. To traverse this distance, I have developed the method of *high-frequency acceleration matching*, which places a high-bandwidth accelerometer at the tip of the remote robot. By inverting a good dynamic model of the haptic interface's acceleration output capabilities, the updated controller can recreate the measured slave-tip accelerations at the user's hand, overlaying this high-frequency haptic feedback channel with traditional position-position control.

High-frequency hand accelerations are similarly difficult to create with closed-loop position feedback in virtual environments, as the renderable spring stiffness is severely limited by stability requirements. Even holding the motor perfectly still, as with a mechanical brake, would not create realistic contact transients because the frequency and duration of the contact response would then be governed by the haptic interface's structural dynamics rather than by the virtual object's properties. To overcome this limitation, I again apply the principles of *high-frequency acceleration matching* to augment the closed-loop controller with an auxiliary high-frequency feedback channel. Contact accelerations are pre-recorded from a real human-object interaction and then conditioned by the inverse of a dynamic model of the haptic interface's acceleration output capabilities. In this manner, the updated controller can recreate measured hand accelerations, overlaying this open-loop high-frequency haptic feedback channel with traditional position control for short durations after contact.

1.6.2 Stabilizing Strong High-Frequency Haptic Feedback

High-magnitude, high-frequency haptic feedback often causes instability in haptic interfaces, especially during teleoperation. Researchers who use position-force control consistently find that their systems exhibit contact instability when the force feedback gain is increased above a certain level, e.g. [56]. Lowering this gain restores stable operation but often does not provide the operator with sufficient stimulation from the remote environment. To improve this poor performance, I have developed the method of *canceling induced master motion*, in which the movement of the master motor in response to the force feedback signal is

estimated and removed from the position command for the remote robot. By canceling induced master motion, the controller can obtain a good estimate of the user's intentional, low-frequency hand movement, thereby breaking the closed loop that becomes unstable under high gain. When the model of the haptic interface is suitably accurate, most of the destabilizing high-frequency motion is canceled, and stable hard contact is possible at much higher haptic feedback levels.

1.7 Thesis Overview

This thesis presents the motivations and methodologies for characterizing and controlling the high-frequency dynamics of haptic interfaces. After the high-level introduction provided by Chapter 1, a historical and technical framework for this research is developed in Chapter 2. The most common feedback methods used with impedance-type haptic interfaces are formally defined, and a variety of existing strategies for improving performance are discussed.

Chapter 3 builds the foundation of this thesis by describing typical haptic interface dynamics and presenting a unified approach to characterization. Common haptic interface models are shown to be poor predictors of system behavior for the high-frequency inputs required to render hard contact. Previous modeling methods are discussed, and their strengths are incorporated into my new two-step approach of comprehensive evaluation and successive isolation. These methods are described in detail and then demonstrated on various haptic interfaces, with particular attention paid to the influence of user grip force. The models developed in this chapter are used extensively in the open-loop model-based control adjustments presented in the following two chapters.

The strategy of high-frequency acceleration matching is explained and demonstrated in Chapter 4 as a method for creating realistic fingertip accelerations during remote and virtual haptic interactions. Previous efforts in this domain are summarized, and the relationship between commanded motor current and fingertip acceleration is highlighted as the fundamental influence on the feel of such feedback. Open-loop model inversion can be used to compensate for these intervening dynamics, enabling the output of accurate high-frequency hand accelerations for both teleoperation and virtual environments. The increased realism of this method was verified through a formal human subjects study, which is thoroughly described and discussed.

Chapter 5 presents cancellation of induced master motion for improved stability during high-magnitude haptic feedback. It begins by distinguishing between intended and induced master motion and continues on to derive a simple criterion for contact instability in position-force control. The stabilizing measures used by others are explained and contrasted with my method of estimating and removing unintended master motion from the slave command. Finally, the stabilizing effect of cancellation is demonstrated by experiment on a teleoperation system and is shown to provide a substantial increase in force-feedback gain.

The thesis concludes in Chapter 6 by summarizing the most important insights and contributions of the presented work. Recommended steps for the mechanical, electrical, and computational design of haptic interfaces are then laid out as a guide to providing more realistic feedback at hard contact. The final section suggests extensions and broader applications for this work in the hope that it might improve the capabilities of robotic surgery systems, educational simulators, and other important efforts in haptic interfaces.

Chapter 2

Background

For the vast majority of history, humans have been able to touch only the objects within arm's reach. The development of mechanical tools enabled people to interact with objects beyond this zone, in the way that a rake allows you to rearrange leaves on the ground without needing to stoop over. Mechanical tools have also facilitated the manipulation of hazardous environments, in the way that long metal tongs can be used to adjust the burning logs in a fireplace. Such devices extend the user's powers of actuation and perception but are still limited by physical proximity.

Over the last sixty years, electromechanical tools have been developed that allow people to touch objects that are located far away or that exist only as computerized models. The interfaces that facilitate this connection are controlled in a variety of ways, as explained in Sections 2.1 and 2.2. Haptic feedback has been shown to improve the operator's efficiency and perceived realism of these remote and simulated interactions, as discussed in Section 2.3. Taken as a whole, this chapter aims to provide a historical and technical context for the contributions of this thesis.

2.1 Teleoperation Control Techniques

Teleoperation dates back to nuclear research by Raymond C. Goertz in the 1940s and 1950s, driven by a need for humans to handle radioactive material from behind shielded walls. With the earliest systems, the operator controlled the motion of the manipulator through an array of on-off-on switches, for example flipping a lever to the left to begin counter-clockwise wrist rotation [45]. Providing "neither force indication nor force reflection (feel),"

these manipulators were "slow and somewhat awkward to operate," leading Goertz to build pairs of master-slave robots whose motions were mechanically linked together via gears and cables, like an extended tool [45, 46]. These new systems allowed the operator to use natural hand motions to move the master mechanism and thus the slave, feeling forces and vibrations from the remote interaction through the connecting structure.

Though a significant improvement over a switch-based interface, a mechanical connection mandates close proximity of operator and environment, limits the system to kinematically identical devices, and requires the user's arm to generate all interaction forces. As early as 1952, Goertz recognized the promise of electrically coupled master-slave manipulators and began an investigation of bilateral force-reflecting positional servos [47]. His concluding observations that a telerobotic system's "stability problem is quite complicated" and "more research and development is needed to improve the performance" aptly foreshadowed the last fifty-four years of work, as engineers and scientists have juggled the dual objectives of stability and haptic feel, searching for ways to bring the user closer to the remote environment.

2.1.1 Position-Position Control

After investigating several alternatives, Goertz decided to connect his electronic master and slave manipulators via bilateral proportional-derivative (PD) control [48], which is now also known as position-position control. Measuring position and velocity errors, this controller creates a virtual spring and damper between the motors of the two devices, pulling them together in an attempt to emulate the direct mechanical connection of earlier systems. Goertz chose this architecture for its simplicity and stability, the same reasons it was adopted in the first computer-controlled master-slave system in 1976 [149], and why it is the most commonly used control method in today's telerobotics.

Throughout the following discussion, I examine the behavior of a teleoperation system by viewing it as a long chain connecting the user to the environment, including the hand, handle, master motor, controller, slave motor, and end effector. This dynamic chain is depicted via lumped parameter elements in Figure 2.1. The user's influence on the system is approximated as a position set-point attached to the master handle through a spring and damper that represent the arm muscles and skin. The environment with which the slave tip is interacting is also portrayed as a spring and damper, along with a frictional contact. For clarity, only one member is depicted between the endpoint and motor of both master

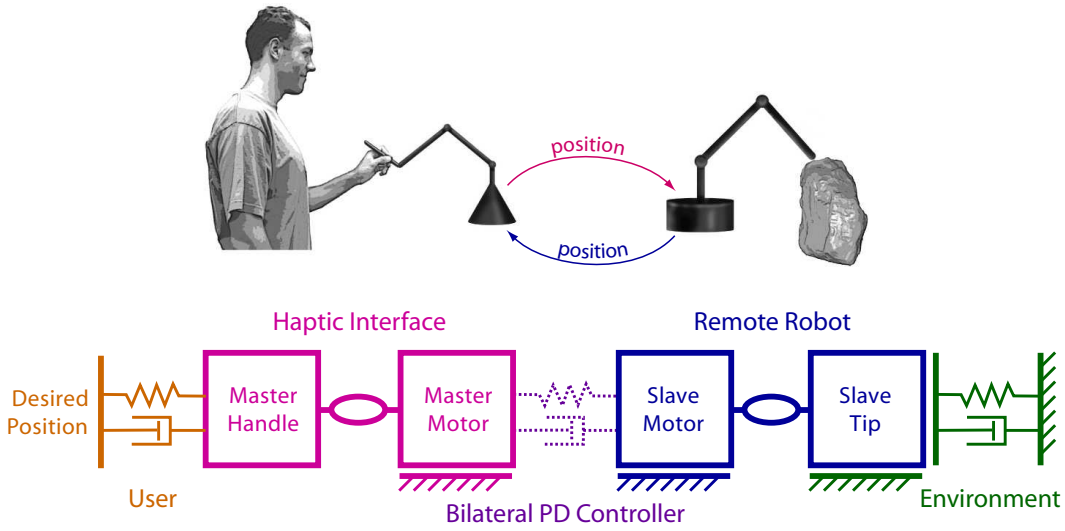

Figure 2.1: Illustration and single-axis lumped-parameter model of a telerobotic system under bilateral PD control.

and slave mechanisms, though these connections can include several spring-like, dissipative, and inertial elements. Based on this treatment, the position-position controller chosen by Goertz and many others can be represented by the dotted spring and damper that connect master and slave motors in Figure 2.1.

Despite its attractive two-sensor simplicity, practical limitations prevent position-position control from joining master and slave tightly enough for realistic haptic feedback. Even without the influences of a user or the environment, the two robots form a closed-loop system that has its own dynamics. Sensor discretization, actuator dynamics, time delay, and structural compliance all compromise the stability of this loop at high gain, limiting closed-loop bandwidth to about five to 20 hertz, e.g. [16]. Although this control methodology adequately transmits intentional hand motions and quasi-static forces, it cannot feed back the broad spectrum of forces that humans can sense during manual interactions, leaving stiff objects feeling soft and undefined.

2.1.2 Position-Force Control

The alternative strategy of position-force control seeks to provide accurate feedback by explicitly measuring the force of contact between the slave and the environment. As illustrated in Figure 2.2, the slave is commanded to follow the measured position of the master

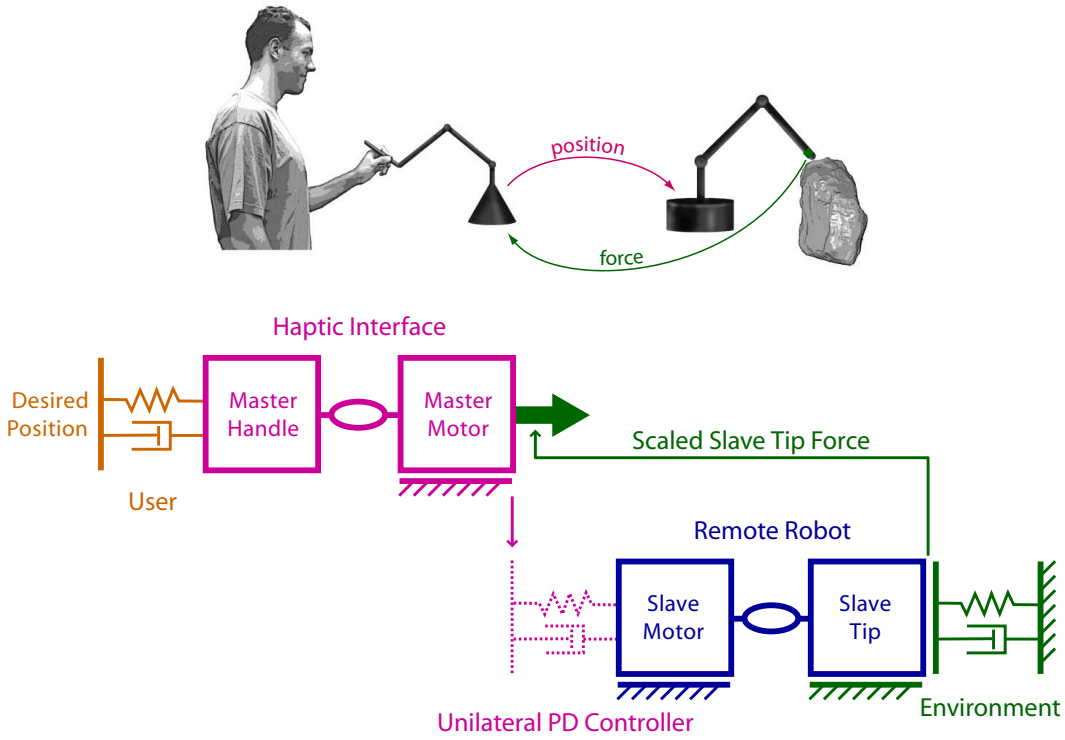

Figure 2.2: Illustration and single-axis lumped-parameter model of a telerobotic system under position-force control.

mechanism via proportional, derivative, and sometimes integral feedback. Contact forces at the slave's end effector are measured with a dedicated sensor and simultaneously displayed via the master's motors. This haptic feedback is transmitted to the user's hand via the structure of the interface so that the user can feel the effects of the slave's environmental interaction.

Although it provides a more direct path from environment to user and hides the slave's friction and inertia, this architecture suffers from contact instability, as feedback forces trigger master motions that excite further contact forces. Forces must typically be attenuated to prevent closed-loop feedback from driving the system unstable [88], again trading off stability and performance. Additionally, all high-frequency feedback forces are distorted by the dynamic chain between the master mechanism's motors and the human's hand. This chain often includes several lightly damped structural resonances, which distort the user's haptic perception of the remote environment and interfere with material and texture identification. Despite its drawbacks, position-force control does allow the user to receive some

high-frequency feedback during contact with remote objects, reducing completion time and the application of extraneous forces in tasks such as peg-in-hole insertion [56].

2.1.3 Two-Port Networks and Time Delay

Another important perspective for understanding and adjusting closed-loop telerobotic behavior is that of the two-port network, which treats the connection between user and environment as a series of energy exchanges [54,123]. This model enables interpretation of each system element as a two-port impedance and inspires a passivity approach, where energy dissipation can be used to determine stability. The performance of such two-port networks has been studied extensively, specifically under the condition of time-delayed transmissions, which easily destabilize traditional controllers [88]. Lawrence argued that a transparent connection requires the communication of both position and force information from master to slave as well as from slave to master. This more complicated control structure is called the four-channel architecture, and it can achieve greater stability and transparency than traditional approaches when properly tuned.

Viewed as a two-port network, a telerobotic system will be stable when all of its elements are passive, i.e. when they do not generate energy. Communications time delay has been shown to violate passivity, which was initially addressed by Anderson and Spong through the use of scattering operators [3] and later handled by Niemeyer and Slotine's strategy of wave variable encoding [107, 109]. Although passive transmissions have been noted to reduce the perceived stiffness of the environment [4, 87], appropriate design choices can improve the feel of such systems, providing consistent, predictable behavior in addition to guaranteed stability under arbitrary time delay [110, 112]. The user's experience of free-space motion and hard contact can be further improved by automatic online tuning of the wave impedance [141], and even better feedback can be provided to the user by adding a slave force sensor and appropriate wave filters [142]. This focus on the user's ability to adequately feel the environment has emerged as a common yet difficult-to-achieve goal of modern teleoperation.

2.1.4 Vibrotactile Feedback

Contact realism has also been improved by displaying vibrations via the master device in teleoperation. Kontarinis and Howe overlaid accelerations measured at the slave's end

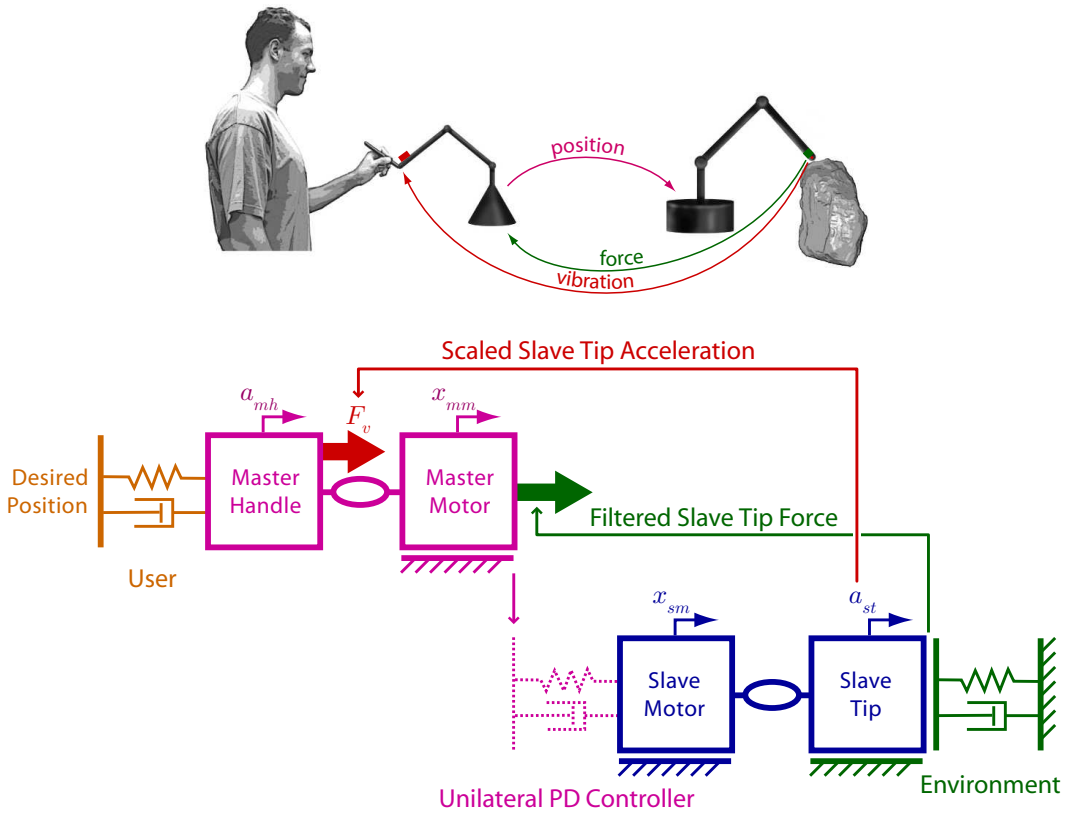

Figure 2.3: Illustration and single-axis lumped-parameter model of a telerobotic system under position-force control with additional vibrotactile feedback.

effector with traditional position-position feedback via a supplementary voice coil actuator mounted near the user's fingertips [76], as illustrated in Figure 2.3. Careful placement of this vibrating element with the purposeful use of an intervening compliance, combined with the small magnitude and high frequency of the displayed signals, allowed them to decouple the feedback and command paths, preventing closed-loop instability. User tests indicated that this hybrid feedback strategy increased user performance in inspection, puncturing and peg-in-slot tasks, and later work improved the vibration actuator for better rendering of contact transients [30]. The merits and drawbacks of this approach are analyzed further in Section 4.1.1.

2.1.5 Sensory Substitution

To avoid closed-loop instability, other researchers have tried to provide force cues via the means of sensory substitution. Hawkes measured accelerations at the slave robot fingertips, playing the signal to the operator as a sound via headphones rather than with a master-mounted actuator [59]. In another approach, Massimino and Sheridan presented low-frequency force information to users of a telerobotic system using audio as well as vibrational information, finding an improvement over visual feedback alone [99, 100]; this control method falls under sensory substitution because the auditory and vibrational waveforms mapped frequency to contact location rather than to the dynamics of the impact itself.

In another instance of sensory substitution, Debus et al. used fixed-frequency vibrations from two secondary handle-mounted actuators to convey whether the operator of a telerobotic contour-following task was contacting the environment with an appropriate amount of force [29]; one side of the custom-designed handle would vibrate when the applied force was too low, and the other side would vibrate when it was too high. The addition of these vibration signals allowed users to modulate applied force more accurately than with visual feedback alone, and it also improved performance when overlaid with standard position-position control. Griffin et al. employed similar strategies in a teleoperated object-handling task, finding that auditory cues provided during the application of extreme grasp force allowed subjects to reduce the total force applied to a transported object [51]. All of these methodologies circumvent the closed-loop stability problems posed by high-frequency force feedback, but they therefore cannot capitalize on the user's manual experience, creating inherently unnatural, cognitively intense interfaces.

2.2 Virtual Environment Control Techniques

While the haptic feedback provided during teleoperation can stem directly from the remote robot's interaction with its surroundings, the system controller for a virtual haptic interaction must determine its output based on an environment model. The earliest interactive virtual environments were computerized flight simulators, which were developed from earlier electromechanical models in the 1970s as an efficient method to safely train commercial and military pilots [119]. These systems typically consisted of a mock cockpit surrounded by visual displays and mounted on a moving platform. In modern systems, a sophisticated

computer monitors the operator's control inputs and simulates the motion of the airplane, shaking and tipping the cockpit appropriately. Whole-body situational simulators are now used in a variety of applications, including flying, driving, and even locomotion [21, 27, 126]. While these dynamic systems differ from the physical object interactions necessary for hand-guided exploration and manipulation, their continued success speaks well to the potential of virtual environments for training other types of experts in complex physical tasks.

Haptic simulation of the interactions between a physical object and a user's hand (or a hand-held tool) began in the mid-1990s with the development of the Phantom haptic interface [98]. This impedance-type device has since been used in a wide variety of applications, including assembly path planning, deformable object simulation, computer games, medical and surgical simulation, and three-dimensional design [132]. Despite the breadth of the virtual environments that have been developed, the control methodologies that underlie their operation are very consistent.

2.2.1 Position Feedback

The geometry of the objects in a virtual environment can be defined either implicitly via mathematical functions or explicitly via point locations joined by a polygon mesh. The user's hand movement is measured via the encoders on the haptic interface's motors and mapped into the virtual environment as a point, sphere, or other shape. Collision detection is performed between the user's proxy and the surfaces of the environment to determine the occurrence of contact [158]. When contact takes place, the direction of the surface normal is computed, and the penetration vector $\vec{x}$ along this axis is determined. The classic haptic rendering algorithm for rigid objects [9] applies a force $\vec{F}$ proportional and opposite to the user's penetration into the virtual object,

$$\vec{F}(t) = -k\vec{x}(t) \tag{2.1}$$

where k represents the virtual object stiffness and t is time. The haptic feedback force has zero magnitude when the user is not inter-penetrating a virtual object, since $||\vec{x}|| \equiv 0$ in this case. This algorithm can be visualized as a virtual spring between the motor of the haptic interface and the surface of the virtual object, as illustrated in the lumped parameter model of Figure 2.4.

The basic approach of proportional feedback captures steady-state interaction forces, but

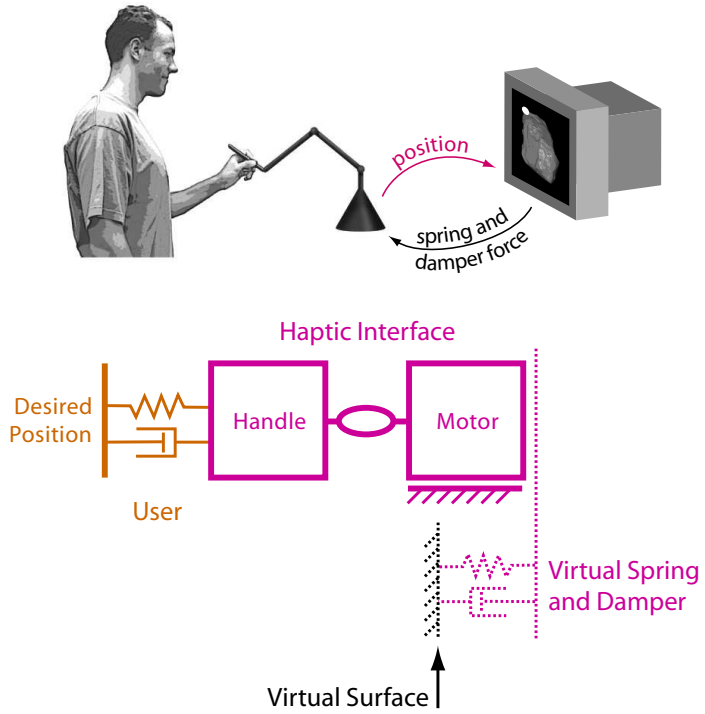

Figure 2.4: Illustration and single-axis lumped parameter model of a user touching a virtual environment through a haptic interface under position and velocity feedback.

it cannot accurately render high-frequency contact dynamics. Fundamentally, this penalty-based method represents a linear control system: the frequency content of the force output is governed by the input signal (i.e. the user's motion) and the system's resonant frequency, $\omega = \sqrt{k/m}$, with m denoting the combined mass of the haptic device and the user's hand. Assuming a typical gain of $k = 1000$ N/m, limited by sensor resolution and computational update rates, this resonance falls no higher than 15 hertz. In contrast, interactions with real materials such as metal and glass commonly exhibit force transients above 1000 hertz [114]. Closed-loop generation of such signals would require feedback gains up to 1,000,000 N/m, roughly a 1000-fold increase over the performance of existing systems; such stiffnesses would necessitate significant improvements in both position resolution and servo rate, as analyzed in [32]. Considerable effort has been expended to understand and overcome these limitations [23, 43], but closed-loop feedback is inherently restricted to smooth, low-frequency forces.

2.2.2 Velocity Feedback

Position-based haptic feedback in virtual environments is doomed to render soft, dull contacts devoid of the information-laden high-frequency transients encountered during real manipulation. As illustrated by the virtual dashpot in parallel with the virtual spring in Figure 2.4, this algorithm can be augmented to include velocity feedback; adding viscous damping to oppose penetration when the user is entering a virtual object improves the crispness of contact. Lawrence et al. explained this effect by distinguishing between the stiffness of virtual walls and their perceptual hardness [90]. They defined the metric of rate-hardness to quantify a virtual wall's ability to display quickly changing forces at contact and found it to be a better indicator of perceived hardness than the wall's stiffness [89]. Unfortunately, adding derivative feedback requires a good velocity estimate, which is seldom available on systems with digital position sensors and uncertain servo-loop timing. In practice, few haptic simulations include virtual damping, which contributes to their soft, springy feel.

2.2.3 Event-Based Feedback

The paradigm of event-based haptics defines an alternative display strategy for creating high-frequency contact accelerations [64, 79]. Rather than trying to generate adequate contact transients using closed-loop position and/or velocity feedback, this method uses discrete event triggers to begin playback of pre-computed force histories. An impact event is triggered when the stylus moves through the surface of a virtual object. The entire transient signal is computed and then overlaid with traditional proportional feedback for its short duration, as depicted in Figure 2.5. With tapping, the contact state is latched during transient output, preventing multiple event detections. Relying on open-loop display somewhat relaxes the need for high-gain closed-loop feedback, and event-based haptics has the additional advantage that brief transients can utilize higher peak currents than steady-state feedback without violating the device's thermal restrictions.

Two main classes of parametric contact transients have been previously investigated and are discussed below: pulses and decaying sinusoids. Regardless of the method used to predetermine the transient, its output remains deterministic up to the user's tactile reaction time of approximately 140 ms [92] and hence does not require continual sensor feedback or additional online computation. Several other event-based methods have been developed

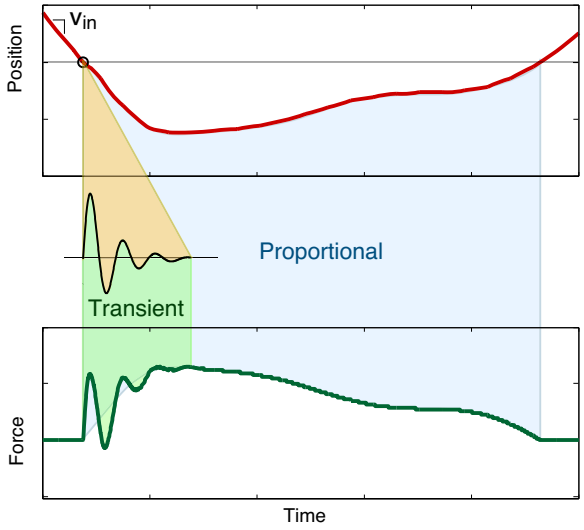

Figure 2.5: Continuous proportional feedback plus an event-based transient at contact.

for augmenting proportional feedback with texture and friction, perturbing the force vector with a variety of different signals for a more realistic feel [41, 96, 104].

Pulse Transients

Several researchers have found that adding pulse-shaped forces to oppose the haptic interface's motion at contact creates sharp force changes and reduces virtual object penetration [24, 64, 128]. Bringing the hand to a stop requires the force transient to be carefully matched to the user's incoming momentum. Under the assumption of constant mass, the system's linear momentum at contact is proportional to the hand's incoming velocity v_{in}. Parametric transient signal magnitude is thus usually scaled by this measurable quantity. The transient's other parameters, such as pulse magnitude per unit velocity and pulse width, are typically tuned by hand to achieve an appropriate feel, though studies have not yet been conducted to evaluate their perceptual impact.

Decaying Sinusoid Transients

Decaying sinusoid transients have also been used to great effect in event-based haptics. Okamura et al. demonstrated significant improvements in perceived realism and task completion during virtual tapping, stroking, and puncturing tasks that included the display of velocity-scaled decaying sinusoid transients at key moments [114]. This study shows that

a combination of vibration and force feedback, displayed simultaneously via the master mechanism's motors, heightens perception capabilities for material discrimination tasks in virtual environments. Unfortunately, signal parameters such as frequency, magnitude, and duration must again be hand-tuned for an appropriate feel, though extensive user testing can also be employed [113].

2.3 The Benefits of Haptic Feedback

As our understanding of stability and control has grown, many researchers have turned their attention to the human experience of haptic feedback. Although tasks can usually be performed without it, force feedback has been shown to reduce completion time, error incidence, and excessive force application during general telerobotic manipulations such as peg-in-hole insertion [56]. Haptic feedback was shown to be especially useful in situations where remote vision fidelity was reduced due to low frame rate or subtended visual angle [100]. Furthermore, while visual feedback can convey sufficient information for the exploration of soft environments, surface deformations of hard objects are too small and abrupt to be seen. The same argument applies to the exploration of edges, small surface features, and other discontinuities; in all such cases, visual cues are of limited value, and the sense of touch becomes the dominant source of information [75].

Interactions with a virtual environment also feel more natural with haptic feedback. Completion time of a rapid tapping task was reduced by adding contact feedback to the virtual targets [5], and simple vibratory feedback during drag-and-drop computer tasks has been shown to ease the user's cognitive workload, determined via both quantitative and qualitative measures [150]. Over the last ten years, a variety of haptic feedback algorithms have been shown to improve user interactions with virtual objects [14, 139], and specific benefits have been identified in tasks such as surgical training [8, 103, 145].

Many experts believe that the intricate tactile nature of surgery, from suture tying to palpation, will be best served by telerobotic systems that provide high-fidelity feedback [19, 134, 144]. One early effort found that position-position-based signals (which include inertial and frictional forces from the slave) fatigued the operator during mock surgery and did not allow perception of soft tissue contact [94]. Indeed, the present generation of commercial telerobotic surgery systems has succeeded partially because soft tissue manipulation can be performed without haptic feedback if high-resolution stereoscopic vision capabilities are

provided. More recent investigations have documented haptic feedback's ability to reduce the excessive application of force, both in blunt dissection [151] and in suturing [74], though further improvements will be necessary to match the dexterity of direct manipulation.

The positive influence of force feedback on manual tasks can be understood by dividing manipulation into a low-frequency power band and a high-frequency information band, as suggested by Daniel and McAree [28]. As discussed in Section 1.1, human manipulation is inherently asymmetric; low-frequency motions elicit high-frequency dynamic responses from the environment. These sudden accelerations, with frequency content generally on the order of several hundred hertz, strongly stimulate the Pacinian corpuscles in the human hand and fingertips [154]. Such tool-mediated vibrations provide rich information about an object's properties, including geometry, texture, and material composition; for example, when exploring real items with a stylus, LaMotte's subjects achieved significantly better levels of discrimination via tapping as opposed to pressing, adeptly inferring softness from the high-frequency dynamics elicited by impact [86]. Clearly, high-frequency accelerations provide important manipulation cues that should be rendered by high-fidelity haptic systems.

Chapter 3

Haptic Interface Dynamics

A haptic interface plays the important role of connecting the user to the controller during interactions with remote and virtual objects. As noted in Section 1.2.2, this research is concerned specifically with actuated impedance-type interfaces due to their excellent free-space characteristics and their widespread use in a variety of applications. During an interaction, the controller of an impedance-type device must measure the user's hand motion and apply an appropriate force in response. Impedance-type haptic interfaces vary in design, but they usually include a series of dynamic elements between the handle and the computer, as described in Section 3.1. Although the common device models discussed in Section 3.2 assume that this connection behaves in a simple manner, the separation between the user's hand and the motor typically results in high-frequency dynamics that are quite complex.

Almost all of the control methods presented in Chapter 2 were developed assuming a simple model for the haptic interface or without a model at all. This design choice limits the usefulness of such approaches and provides an opportunity for substantial improvement. As such, I have sought to characterize the internal dynamics of typical haptic interfaces, building models that accurately predict system behavior, especially at high frequency. As presented in Chapters 4 and 5, these models can then be used to improve both the feel and the stability of a haptic system. Some previous efforts in device identification are examined in Section 3.3, and the first core element to my characterization approach is explained in Section 3.4. Comprehensive evaluation is extended to the time-varying parameter of user grip force in Section 3.4.6, and Section 3.5 presents a thorough

methodology for building a physically based model of the entire dynamic chain. As summarized in Section 3.6, the main goals of this chapter are to convey the nonideality of internal haptic interface dynamics and to provide a reliable set of methods for their identification.

3.1 Elements of Impedance-Type Haptic Interfaces

Haptic interfaces usually provide two or three degrees of freedom in position, sensing the user's motion and applying feedback forces within this workspace. Many devices permit changes in the orientation of the end effector; these rotational degrees of freedom can be unsensed, sensed but not actuated, or sensed and actuated. The remainder of this work will focus on translation rather than orientation, though the developed methods can be applied to either.

Figure 3.1 illustrates the chain of elements typically present in each axis of a haptic interface. For clarity, the illustration depicts a device with a single degree of freedom, but typical systems combine several degrees of freedom in parallel or series to allow unrestricted translation and/or orientation. Although differences exist, individual position axes of most mechanisms can be represented by such an arrangement. For the duration of this thesis, the terms "haptic interface" and "master" will be used interchangeably to represent all electrical and mechanical elements depicted in Figure 3.1, extending from the amplifier and encoder to the handle.

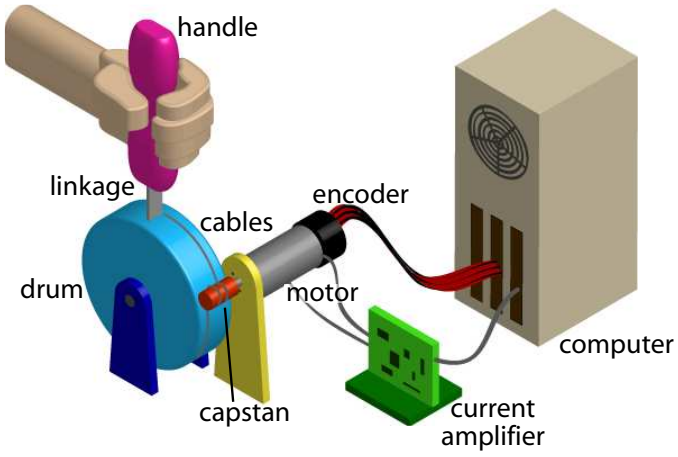

Figure 3.1: One axis of an impedance-type haptic interface.

3.1.1 Computer

The haptic interface's controller typically runs on a real-time enabled computer at a fixed servo rate, which is often one kilohertz. Haptic feedback forces are converted to a desired current for each DC motor using the mechanism's Jacobian-transpose matrix, the gear ratio of each joint, and each motor's torque constant. These current commands are communicated to a set of self-contained amplifiers through a digital-to-analog converter (DAC) that often resides on a control card on the computer's ISA or PCI bus. Some amplifiers also accept digital commands, communicated over a parallel or serial connection.

3.1.2 Current Amplifier

Each amplifier is connected to one motor, and it attempts to regulate motor current via pulse-width modulation (PWM) or linear control techniques. PWM amplifiers are presently somewhat more common in haptics due to their widespread use in industrial robotics where their lower power consumption is important. Unfortunately, PWM amplifiers generate significant high-frequency electrical noise at their switching rate and its harmonics, which can contaminate analog sensor lines. Additionally, PWM amplifiers are often tuned by the manufacturer to a low bandwidth, often on the order of 100 hertz, which is adequate for industrial applications but must be increased for high-frequency haptic interaction. If their additional power consumption can be tolerated, linear amplifiers are generally preferable, as they can provide very clean, high-bandwidth current output without interfering with the system's sensing requirements.

3.1.3 Motor

Haptic interfaces typically use small, brushed DC motors such as those available from Maxon Precision Motors, Inc. [101], as they provide very smooth torque generation and have high power to weight ratios. Current flowing through the motor creates a torque on the motor shaft, to which a small capstan is attached. The relationship between the motor current, i_m, and the applied motor torque, τ_m, is governed by the motor's torque constant, k_t, as follows:

$$\tau_m = k_t \, i_m. \tag{3.1}$$

The torque constant for a motor can be obtained from the manufacturer's data sheet and can also be calibrated using an ammeter and a torque sensor. When used in a haptic

interface, the net torque output of the motor will be diminished by the friction present at the motor shaft, so low-friction motors and bearings are desirable.

Another important set of characteristics for a DC motor are its heat dissipation capabilities and its internal temperature limit. The flow of current through the motor coils produces heat that raises the temperature of the rotor. A motor that is driven with high levels of current that cause it to exceed its internal thermal limit will burn out and need to be replaced. This phenomenon is often viewed as setting a maximum steady-state current that a motor can sustain indefinitely. Most interface designers conservatively choose to operate under this limit at all times, though more sophisticated thermal monitoring schemes can also be employed [38]. As another consideration, a rise in motor temperature increases the motor's electrical resistance and therefore reduces its electrical efficiency.

3.1.4 Encoder

Motion of the haptic interface is usually sensed with an optical encoder attached to the motor shaft. This type of sensor provides two digital output lines, often denoted A and B, that stem from two optical sensors in the encoder. These two sensors are pointed at a reflective disk that has many thin radial lines cut out of it or painted onto it; this disk rotates with the motor shaft. Each sensor reads high and low as lines pass before it, and their locations are chosen to place the signals 90° out of phase from one another. A quadrature decoder chip, which is usually located on a control card on the computer's ISA or PCI bus, observes the output of these two sensors to determine the present angular position of the motor shaft. The output of the quadrature decoding is an signed integer that designates the number of ticks the shaft has rotated away from an arbitrary zero location. Each tick represents one quarter of one line on the disk; haptic interface programmers can determine the number of lines in an encoder either from the manufacturer's data sheet or from calibration. This information enables computation of the encoder resolution, Δ, which is measured in radians per tick and is calculated as

$$\Delta = \frac{4\,n}{2\,\pi}, \tag{3.2}$$

where n is the number of encoder lines per revolution, commonly between 500 and 2500.

Once the resolution of the encoder is known, the digital output from the quadrature

decoder chip can be transformed into a quantized motor angle reading, θ_m, as follows:

$$\theta_m = \Delta(q - q_{zero}), \tag{3.3}$$

with q standing for the present quadrature output and q_{zero} being a calibration value. This zero offset must be determined every time the system is initialized, often by recording the quadrature readings at a certain known position in the device's workspace.

3.1.5 Cables

Thin stranded cables couple motion of the motor's capstan to that of a larger drum. Cable drives are used instead of belts or gears to enable very smooth, efficient motion of the device [130]; as discussed in Section 1.1, the human hand is very sensitive to high-frequency vibrations, so non-vibratory transmission elements must be used to maintain the realism of free-space motion. When pre-tensioned, the low-stretch, highly-stranded cables available from manufacturers like Sava Industries, Inc. [131], provide a zero-backlash connection between capstan and drum, which is important for ensuring a close coupling between the user's hand and the motor.

3.1.6 Drum

The drum diameter, d_d, is typically ten to twenty times as large as the capstan diameter, d_c, providing the unitless gear ratio, ρ, as follows:

$$\rho = \frac{d_d}{d_c}. \tag{3.4}$$

Assuming that they are perfectly inextensible, the cables couple the motion of the capstan and drum together by this gear ratio with the following two equations:

$$\tau_d = \rho\, \tau_m \tag{3.5}$$

$$\omega_m = \rho\, \omega_d, \tag{3.6}$$

where τ is a torque, ω is an angular velocity, and the subscripts d and m denote drum and motor respectively. The cable drive thus serves the dual objectives of amplifying the motor's torque to enable stronger haptic feedback and amplifying the drum's motion to enable higher

resolution position measurement. The primary disadvantage of a high gear ratio is that it also increases the effect of the motor's rotational inertia and rotational friction at the user's hand, relationships that go with ρ^2 and ρ respectively. Device designers typically balance the four objectives of torque amplification, motion amplification, inertia minimization, and friction minimization to select an appropriate gear ratio.

3.1.7 Linkage and Handle

The drum is attached to the endpoint of the device through a mechanical linkage, and the user holds a handle, stylus, or thimble at the endpoint. The distance between the rotational axis of the drum and the point of user-handle contact is defined to be h, and it relates the translation of the user's hand to the rotation of the drum. If the linkage were perfectly stiff, the coupling relationships would be

$$F_f = \frac{\tau_d}{h} = \frac{\rho\, \tau_m}{h} \tag{3.7}$$

$$\omega_m = \frac{\rho\, v_h}{h}, \tag{3.8}$$

where F_f is the haptic feedback force applied to the human, and v_h is the translational velocity of his or her hand. Such devices are designed to have low friction and low inertia so that the user can easily move them by hand and so that the applied haptic feedback is more salient than the forces resulting from the natural dynamics of the device.

3.1.8 User

The final element that affects the behavior of a haptic interface is the user. When an individual grasps the handle at the end of the system's long dynamic chain, he or she gains the ability to physically affect its motion and to be affected by it in turn. The skin and muscles of the human hand are somewhat compliant and dissipative, and the flesh and bones of the hand have significant mass [53, 120]. While all of the other elements in the master's dynamic chain are generally time invariant, with the possible exception of configuration dependence, different users consistently possess unique dynamic characteristics. Additionally, the dynamic response of each user can vary over time with changes in grasp configuration and the co-contraction of various muscle groups.

In his work on position-force control in 1988, Hannaford insightfully observed that "unless human operator impedance is taken into account in the real time operation of the control laws, optimum performance must be sacrificed to guarantee safe and stable operation" [55]. He suggested three methods for implementing this strategy, including allowing the user to tune controller parameters during operation, installing a one-bit contact sensor to detect sufficient grip strength, and continuously estimating operator impedance from the relationship between applied force feedback and measured master position. My work naturally continues these tactics by seeking to establish a known correlation between grip force and user impedance, as will be discussed in Section 3.4.6.

3.2 Common Models

Despite the complexity of their many connected elements, haptic interfaces are most often modeled as ideal mechanisms. The discretization of the computerized servo loop is generally neglected, as are any possible variations or delays in the controller's execution time. Current commands and motor position readings are treated as continuously varying signals, ignoring their respective quantizations. Current amplifiers are assumed to have infinite bandwidth and perfect disturbance rejection, and current saturation is seldom modeled. Beyond supposing that these electro-mechanical interface features will not affect the behavior of the system, researchers also often make strong assumptions about the mechanical behavior of the haptic interface, employing two common models that are derived from fundamental physical principles.

3.2.1 Simple Mass

The first model frequently used to represent the behavior of a haptic interface is a simple mass, as depicted in Figure 3.2. This model assumes that the cables, linkage, and handhandle interface all behave as perfectly rigid members, allowing their inertias and the inertias

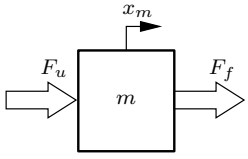

Figure 3.2: Simple mass model of a haptic interface.

of the motor, drum, handle, and hand to be lumped into a single effective mass, m. This model also neglects any friction that may be present at the motor or drum, as haptic interfaces typically exhibit little friction to ground. These assumptions enable the use of (3.5) to (3.8), implying that the system is capable of perfect force production at the user's hand and perfect measurement of the user's hand motion. The device's encoder reading is taken to indicate the position of the mass, x_m, and the motor torque is assumed to directly apply the force F_f. The influence of the user is modeled by the force F_u and is treated as a system input, though it is seldom measured explicitly on impedance-type devices.

While its simplicity is very attractive, the common mass model provides a very poor approximation for the behavior of an actual haptic interface when held by a user. The dynamic equation that governs this model's position response to force feedback is

$$\frac{X_m(s)}{F_f(s)} = \frac{1}{ms^2} \tag{3.9}$$

where s is the Laplace operator. As the system's two input forces are collocated and therefore indistinguishable, the response to user force takes the same form. This double integrator has an unbounded response to any net feedback and user force inputs that have a non-zero mean. A small, constant feedback force that is not countered by an equal and opposite user force is thus predicted to move the device to positive or negative infinity; when held by a user, as it is during all haptic interactions, a haptic interface will not behave in such a manner. First, any real device will have some amount of friction, which acts to remove energy from the system and slow the mass down. Second, the user is often attempting to move the device along a certain trajectory and will typically oppose moderate levels of force feedback to maintain his or her intended path. This model provides no method for differentiating between the passive and active aspects of the user, i.e. their natural hand dynamics and their intentional hand motion. It thus becomes impossible to predict the motion of the device without sensing the force applied by the user at the handle.

Researchers who use this model for its simplicity sometimes acknowledge that it struggles to capture the non-inertial properties of the user. They typically set the user force to zero to approximate the situation when a user is not holding the handle and argue that this scenario represents the most challenging dynamic response for keeping the system stable. I argue that constraining the handle in the user's hand actually exposes the nonidealities of the master's dynamics and can make stability far more difficult to attain, especially when

lightly damped internal resonances exist. Furthermore, the objective of a haptic interface is to allow a human to interact with a distant or virtual environment; neglecting the user's influence seems to defeat the purpose of the system and can result in controller designs that fail to perform as well as predicted.

3.2.2 Mass-Spring-Damper

The other model commonly used to predict the dynamic response of haptic interfaces is that of a mass, spring, and damper, as shown in Figure 3.3. Examples of researchers who use this model type include [72,88,123,133]. Like the simple mass, this model also assumes that the master's mechanical transmission elements (cables and linkage) are perfectly rigid, allowing their inertias to be lumped with the inertias of the motor, drum, handle, and hand to form a single mass, m. The user's attachment to the system is modeled as a spring and damper in parallel, connecting the single mass of the haptic interface to the position that the user desires for his or her hand, x_{des}. The stiffness, k_u, and viscous damping, b_u, represent the action of the user's muscles in response to disturbances from the device. Many researchers use second-order models to describe passive joint dynamics [58, 61, 73], and my previous investigations have supported the efficacy of such a user model in haptic interactions [80, 84]. The user's desired position will generally change at no more than a few hertz, but the spring and damper allow the mass to move at frequencies much higher than this when driven by force feedback.

The mass-spring-damper model predicts low-frequency system behavior much more adeptly than the simple mass model by accounting for both the user's intentional and non-intentional influence on the system. The dynamic equation that governs its position response to force feedback is

$$\frac{X_m(s)}{F_f(s)} = \frac{1}{ms^2 + b_u s + k_u}. \tag{3.10}$$

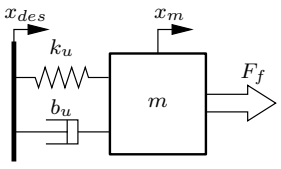

Figure 3.3: Mass-spring-damper model of a haptic interface.

This model yields a finite deflection of F_f/k_u under a constant force ($s = 0$), providing a more tenable prediction of the interface's low-frequency behavior. If its three parameters are well identified, this model type can be used to analyze teleoperation and virtual environment systems that have only low-frequency feedback signals. It generally should not be used to predict the behavior of haptic interfaces subjected to high-frequency force feedback, i.e. signals at frequencies much above the model's resonant frequency of $\sqrt{k_u/m}$.

The mass-spring-damper model of a haptic interface cannot capture the high-frequency behavior of typical systems because it makes several erroneous assumptions. First, effects such as the servo-loop's finite sampling frequency, servo-loop delay, and amplifier dynamics become increasingly significant at higher frequencies. For example, while it is generally acceptable to model an amplifier as a perfect current source at one or ten hertz, amplifiers exhibit low-pass behavior above their rated or tuned bandwidth, attenuating the magnitude and shifting the phase of current requests. These significant effects are not captured by the mass-spring-damper model (nor by the simple spring model), so they have limited efficacy near the servo frequency and near or above the bandwidth of the amplifier.

Second, the assumption of a single lumped mass also breaks down at high frequency. No mechanical connection is infinitely stiff: the cables and the linkage will both exhibit some degree of flexibility, and this behavior significantly changes the system's high-frequency response. Rather than behaving like the lumped mass m at infinite frequency, the system will behave like the motor mass m_m, as the somewhat flexible cables connect it to the rest of the chain. The importance of structural flexibilities is explored by Adams, Moreyra, and Hannaford, who model a haptic interface as an n-length chain of lumped masses connected by springs and dampers [1].

Third, the lumped-mass assumption implies that the user's hand will feel the same accelerations as the motor, which is far from true. The user holds the haptic interface at the handle, which is separated from the motor by several transmission elements and the intermediate mass of the drum. As the user's Pacinian corpuscles are not collocated on the same mass as the actuator, they will undergo high-frequency accelerations that are significantly different from those of the motor. While simple models can capture the low-frequency behavior of a haptic interface, they cannot reproduce these higher-order effects, which become apparent at the high frequencies that need to be displayed to achieve adequate haptic realism for hard contact.

3.3 Prior Work in Haptic Interface Identification

Despite its prominent role of connecting the user to the controller, only a few researchers have sought to characterize the broad-spectrum dynamic response of haptic interfaces. This paucity of investigation can most likely be traced to the low bandwidths of the closed-loop, position-based control methodologies typically employed in teleoperation and virtual environments; developing models that are accurate at high frequency has not been a relevant undertaking until about ten years ago, when a focus on the performance capabilities of haptic interfaces began to emerge. In the context of this thesis, it is beneficial to highlight three previous efforts in haptic interface characterization. These projects were conducted by different sets of researchers for different purposes and are described below.

3.3.1 Structural Deformation Ratio

The first notable approach to system identification was developed by Moreyra and Hannaford in 1998 as a simple method for quantifying the nonideal high-frequency dynamics of a haptic interface [105]. The authors discuss the importance of the dynamic response of the mechanism in creating hard contact interactions that feel "good" to the user. They develop the unitless metric of the structural deformation ratio (SDR) by applying a short impulse in current and measuring the velocity of the motor immediately following the impulse. This measured velocity is divided by the velocity that is predicted by a simple mass model to provide an indication of the flexibility of the mechanism's internal transmission elements; higher ratios indicate that the pulse is accelerating less mass, so the device must have an internal structural flexibility.

Performing the same test for pulses with different durations begins to elucidate the frequency-dependence of the system's dynamics, but it is difficult to build a dynamic model from this metric. Additionally, this method relies on having accurate low-magnitude velocity measurements, which are difficult to obtain with quantized position sensors, and it fails to differentiate between electrical, frictional, and structural effects. Overall, though, this work aptly identifies the importance of high-frequency dynamics in haptic interfaces and provides a straightforward method for calculating the SDR on any device without additional sensors or actuators.

3.3.2 Acceleration Throughput

The second important effort in master identification was published by Okamura, Hage, Cutkosky, and Dennerlein in 2000 and 2001 to motivate the slowing down of recorded high-frequency acceleration transients before display in an event-based virtual environment [113, 115]. This group used a function generator to apply sinusoidal current at individual frequencies to the motors of an impedance-type haptic interface. The output from a tip-mounted accelerometer was observed on an oscilloscope while the experimenter lightly held the handle in place, and the relative magnitude was recorded. By performing tests at a wide range of frequencies, the researchers were able to compute the normalized magnitude response of their system, which exhibited a resonance at 120 hertz, a sharp roll-off until 400 hertz, and various peaks and valleys up to 2000 hertz.

This work by Okamura et al. clearly shows that haptic interfaces do not behave in a manner consistent with a simple mass or a mass-spring-damper at high frequency; both of these models would predict a flat frequency response from motor current to handle acceleration at high frequency, rather than a response that varies strongly with frequency. Unfortunately, the relative phase between input and output at each frequency was not recorded, so matching a dynamic model to the presented data is difficult. Furthermore, obtaining a smooth magnitude plot via single-frequency experiments is laborious and time consuming, though it could be made more efficient through computerized commands and measurements. Finally, it is not clear whether the user's influence during testing was representative of the dynamics that would be in place during a virtual interaction; if they were not, due to a light hold or inconsistent experimenter behavior, then the obtained data may not be an accurate representation of the system's likely performance. While several aspects of this identification project could be improved, the presented work does sufficiently motivate the authors' vibration feedback methodology, and it provides excellent impetus for further investigation of the acceleration output capabilities of haptic interfaces.

3.3.3 Motor Response

The third endeavor at identifying the dynamics of a haptic interface was performed by Çavuşoğlu, Feygin, and Tendick in 2002 [18]. The focus of this work was to understand and improve the capabilities of a stock Phantom to enable high-fidelity applications such as bilateral teleoperation, virtual deformable object portrayal, and frequency-based psychophysics

experiments. After replacing the interface's PWM amplifiers with high-bandwidth linear amplifiers, the researchers commanded single-frequency sinusoids in current and measured the magnitude and phase of the motor's position response. They eliminated the effects of gravity by reorienting the device for tests on each axis, and they did not have a user holding on to the handle during the experiments. Testing elucidated double-integrator behavior on all three axes at frequencies below about 30 hertz. Beyond this frequency, two axes show a significant decrease in response at 50 hertz and a resonance afterwards near 100 hertz, and the third axis shows a minor resonance at 200 hertz; system responses from 200 to 1000 hertz show peaks and valleys in magnitude and a marked decrease in phase.

This research by Çavuşoğlu et al. provides a rich view into the dynamics of the most commonly used haptic interface, showing that good models must include at least two masses in order to capture observed behavior. In some ways, this work is very similar to the acceleration throughput approach described above, though output phase was recorded, and commands and measurements were computerized for increased efficiency. Unfortunately, these tests did not include the influence of a user's hand, presumably to focus on time-invariant device characteristics. When compared with Okamura et al. the primary distinction of this work is that it looked at motor position rather than tip acceleration, a tactic that is somewhat similar to the SDR work discussed above. The experimentally determined transfer functions (magnitude and phase together) obtained by the researchers would be useful in predicting the master's influence on stability in high-bandwidth teleoperation without the influence of a user, but they do not enable us to understand the device's ability to create high-frequency hand accelerations. Valuable work has been conducted in haptic interface identification, providing insights about a variety of mechanisms and methodologies.

3.4 Comprehensive Evaluation

I have developed a set of methods for illuminating the high-frequency dynamics of haptic interfaces and creating models that accurately predict the observed system behavior. Simple second-order models and prior characterization techniques cannot fully capture the dynamic response of the user-master system, so I have created the more sophisticated two-step characterization approach of comprehensive identification and successive isolation. As described in this section, comprehensive evaluation seeks to explore a certain input-output relationship, determining if it behaves in a linear, time-invariant (LTI) fashion and finding

an appropriate LTI model if so. Comprehensive evaluation can also be extended to account for a single time-varying parameter like user grip force, as described in Section 3.4.6. The second step of successive isolation is necessary only for those haptic interfaces that exhibit significant nonlinearities or strongly time-varying behaviors. The detailed dynamic dissection methods of successive isolation are delineated in Section 3.5. Throughout the rest of this chapter, my identification methods are described and then demonstrated on various haptic interfaces.

In all of my identification work, I take the system input to be the current commanded from the amplifier, i_{cmd}, which differs by only a constant scale factor from the force feedback command, F_f. Such a choice allows for identification of the computer's and amplifier's electrical dynamics, which are often significant at high frequency. For each haptic interface, I consider two outputs: the acceleration of the handle, a_h, which is central to user perception of hard contact and texture, and the movement of the motor, x_m, which is utilized by almost all controllers to estimate the user's hand motion. The relationship between the current command and the two outputs depends to some extent on all of the elements in the long dynamic chain between the controller and the user. The objective of the system identification process is to determine the most important effects and include them in the models. I define the term H_{ha} to stand for the relationship from current command to handle acceleration, and I define the expression G_i to represent the relationship from current command to motor motion, which I call induced master motion since it does not stem from the user's volition.

As described in the following sections, comprehensive evaluation applies current commands that stimulate the haptic interface's high-frequency dynamics and measures the resulting handle acceleration and motor movement. Commands and measurements must be executed with very consistent timing, using a fast servo rate of a few kilohertz on the control computer. All testing is performed with a user holding the handle of the interface in order to showcase the dynamic relationships that will be in effect during actual teleoperation or virtual environment exploration. I seek to determine whether linear, time-invariant models, $\hat{H}_{ha}(s)$ and $\hat{G}_i(s)$, can closely approximate the behavior of the haptic interface's complex electro-mechanical system above about ten hertz. Though such models can be obtained in many ways, non-parametric characterization techniques are particularly well suited to this purpose, as they do not assume a model order. These so-called black-box identification tests apply an input signal that is rich in frequency content and observe the frequency content of the resulting output signals.

3.4.1 Experimental Procedure

With the user holding the handle in a comfortable grip, comprehensive evaluation applies a current command that stimulates a range of frequencies in the haptic interface. While a variety of signals can be used for this purpose, this research employs a swept-sine-wave current command because it can be computed easily in real time. The frequency of a swept sine wave varies linearly from low to high or high to low across its duration. At each time step of a test, the controller requests a certain motor current from the amplifier; time steps are denoted by the index k, which ranges from 0 to $N-1$. The equation for the kth element of a discretized swept-sine-wave signal is

$$i_k = i_{amp} \sin(2\pi f_k k T) \qquad k \in [0, 1, \ \ldots \ , N-1], \tag{3.11}$$

where i_{amp} is the amplitude of the swept sine wave, T is the sample time in seconds, N is the number of elements in the sample, and f_k is the kth frequency in hertz, defined to be

$$f_k = f_i + \frac{k}{2(N-1)}(f_f - f_i), \tag{3.12}$$

where f_i and f_f are the initial and final sinusoid frequencies.

Typical values for the low frequency are one to ten hertz, and for the high frequency are 500 to 2000 hertz. These selections are made to span the device's interesting frequency range without going much below the user's active bandwidth of eight hertz or above the Nyquist frequency, which is half of the sampling rate. The input signal spans approximately two seconds, and it is repeated immediately to reduce the effect of start-up transients and to help ensure a smooth transition between the final output values and their values at the start of the second input signal. The input signal and both output signals (handle acceleration and motor position) are recorded for the duration of the test and written to a single computer file that also includes various testing and interface parameters. Sample data from a single trial are shown in Figure 3.4.

Using a swept sine wave makes the identification process much more efficient than the single-frequency sinusoid method. Rather than taking several seconds per frequency, all frequencies are sampled in the same short test. The test is conducted two to four times for each set of input signal parameters, and the sign of the input signal can be reversed for half of the tests to investigate directional symmetry. The user finds it easier to maintain a

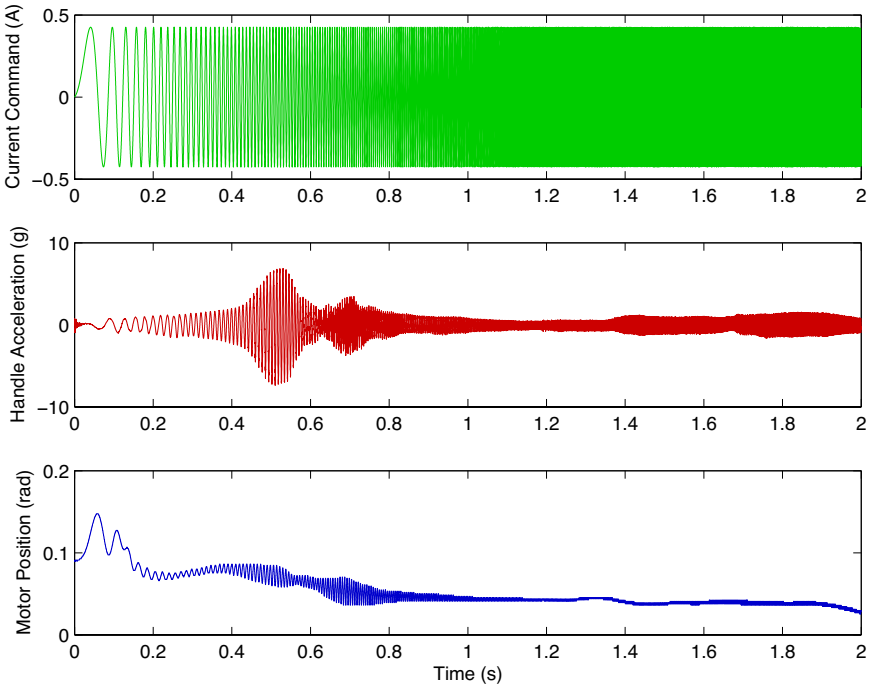

Figure 3.4: Sample swept-sinusoid current command and resulting handle acceleration and motor position.

constant grip force during this low number of trials than over many successive trials, as the hand does not become nearly as fatigued.

3.4.2 ETFE Analysis

An empirical transfer function estimate (ETFE) can be formed from the experimental data by dividing the discrete Fourier transform (DFT) of an output by that of its input. For both handle acceleration and motor position, the system's response to the second of the two swept-sinusoid input signals is first isolated, giving the vectors $\vec{a}_h = [a_{h,0} \ ... \ a_{h,N-1}]^T$ and $\vec{x}_m = [x_{m,0} \ ... \ x_{m,N-1}]^T$. The data are then detrended by subtracting a ramp function to achieve the numeric value of zero at the start and end of the updated vectors $\vec{a}'_h$ and $\vec{x}'_m$, via

$$a'_{h,k} = a_{h,k} - a_{h,0} - \frac{k}{N-1}(a_{h,N-1} - a_{h,0}) \qquad k \in [0, 1, \ ... \ , N-1] \tag{3.13}$$

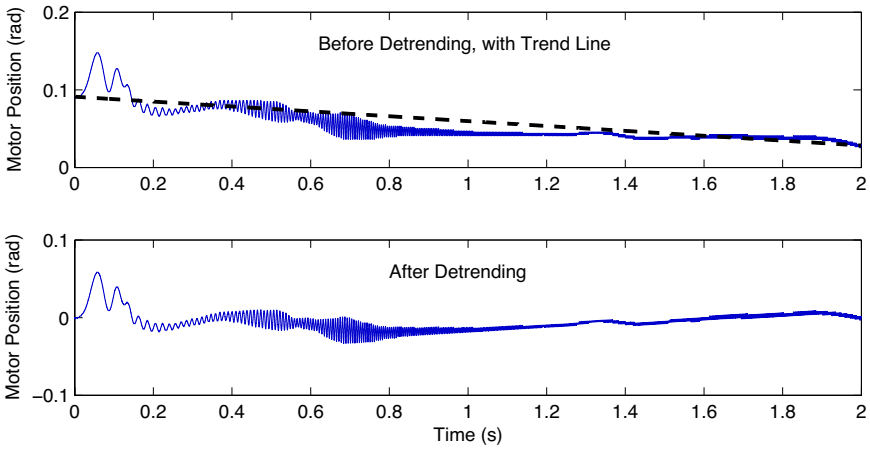

Figure 3.5: Output signal detrending.

and

$$x'_{m,k} = x_{m,k} - x_{m,0} - \frac{k}{N-1}(x_{m,N-1} - x_{m,0}) \qquad k \in [0, 1, \ \ldots \ , N-1]. \tag{3.14}$$

The process of detrending is illustrated for a sample motor position signal in Figure 3.5. If this detrending is not performed, the first and last elements of the vector will not have the same value, and a repeating sequence of the signal will appear to have a sharp jump. This discontinuity will resemble a sinc function in the signal's DFT, which undesirably disrupts the rest of the data. Detrending is especially necessary for the motor position signal because the human hand often drifts many millimeters over the course of a test; subtracting a gentle ramp function from the raw data merely removes this net motion for better viewing in the frequency domain.

Once the handle acceleration and motor position signals have been detrended, the N-length discrete Fourier transforms of the input signal and both output signals are computed. Each transformation yields a complex-valued vector of the same length N as the original data, where the nth element of each DFT vector $\vec{Y}$ is computed by

$$Y_n = \frac{1}{N} \sum_{k=0}^{N-1} y_k e^{-2\pi j n(k+1)/N} \tag{3.15}$$

with $\vec{y}$ being the time-domain vector that is being transformed and j standing for the imaginary unit. The nth element of $\vec{Y}$ represents the magnitude and phase of the signal at the frequency

$$f_n = 2\pi \frac{n}{NT}. \tag{3.16}$$

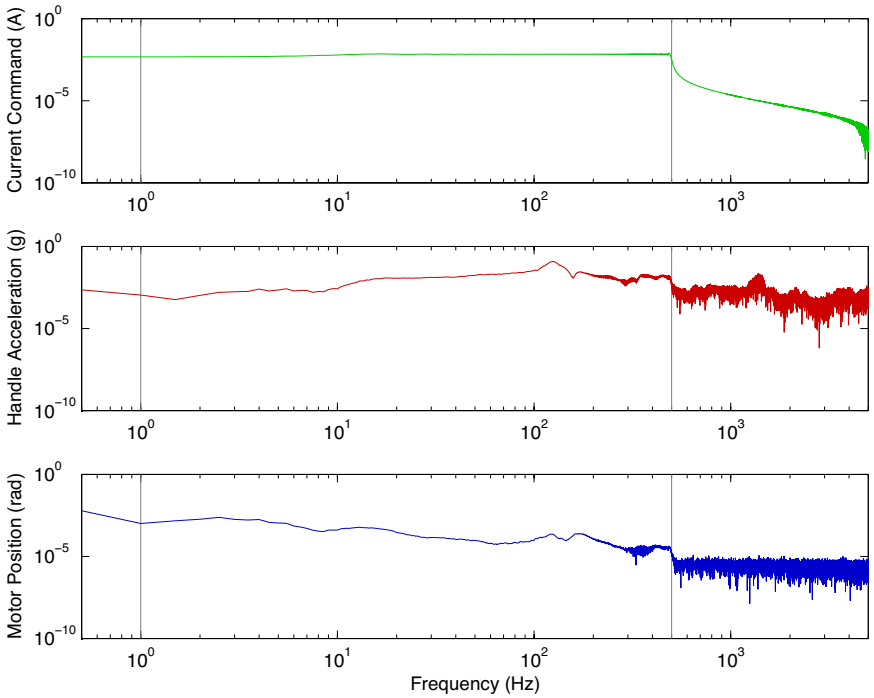

Figure 3.6: Magnitude of the discrete Fourier transform of each signal in Figure 3.4 after detrending.

The second half of the DFT vector contains the complex conjugates of the first half, i.e. $Y_{N-n} = \bar{Y}_n$. The magnitude of the first half of the DFTs of the signals shown in Figure 3.4 are plotted in Figure 3.6 for comparison. Gray lines indicate the starting and ending frequencies of the input signal, i_{cmd}; frequencies outside of this range did not receive significant stimulation and thus should not be considered in the results.

The output DFTs from several tests under the same experimental conditions are averaged in the complex domain to diminish the effect of slight variations between trials. The ratio of the averaged output DFT to the input DFT is computed, and the magnitude and phase of this complex-valued vector are computed. The magnitude and phase vectors are then smoothed using a five-sample boxcar filter and are plotted against the frequency vector $\vec{f}$, as defined in (3.16). The results then comprise an experimentally determined Bode plot, providing insight into the system's broad-spectrum dynamic behavior. Sample ETFEs for handle acceleration and motor position are shown in Figure 3.7.

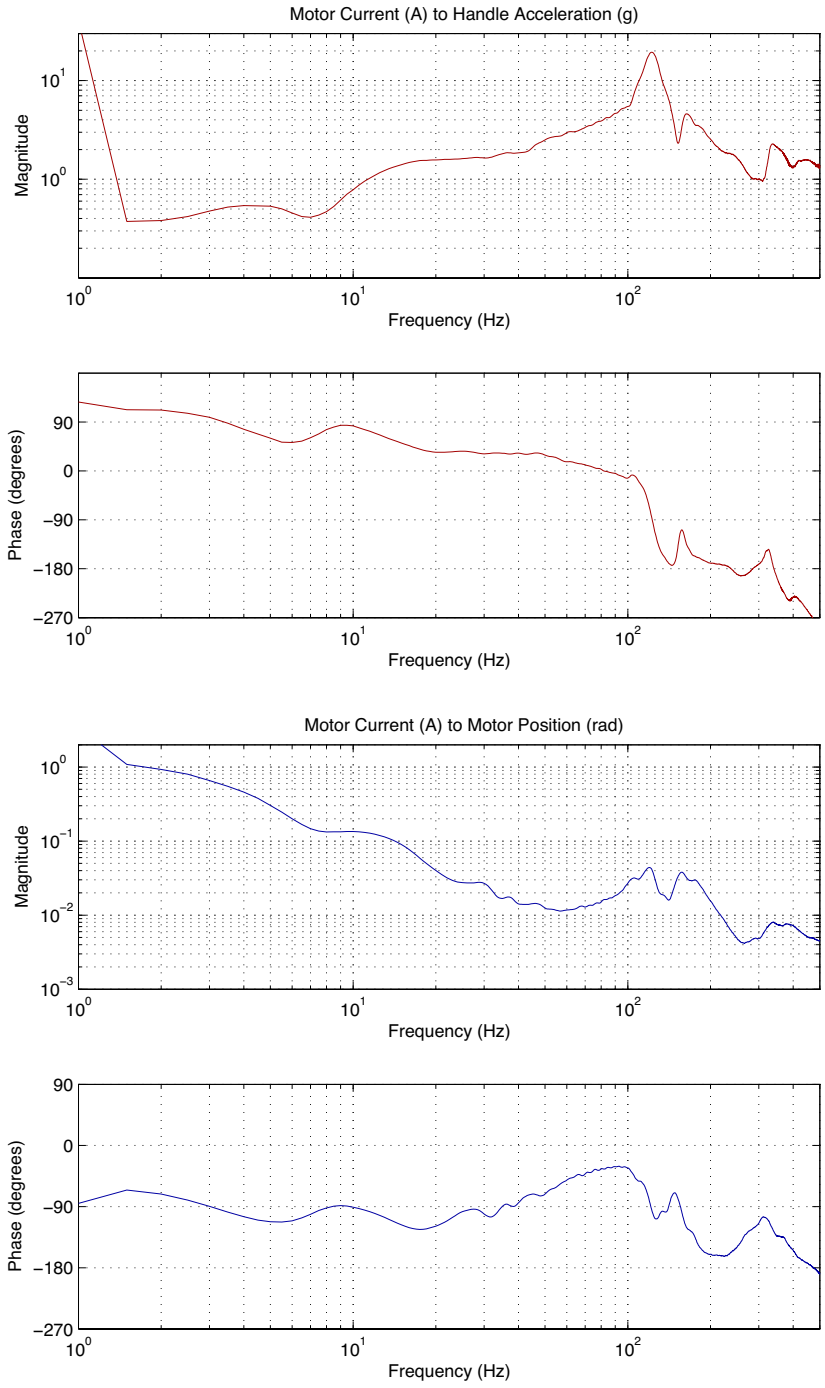

Figure 3.7: Empirical transfer function estimates for handle acceleration and motor position.

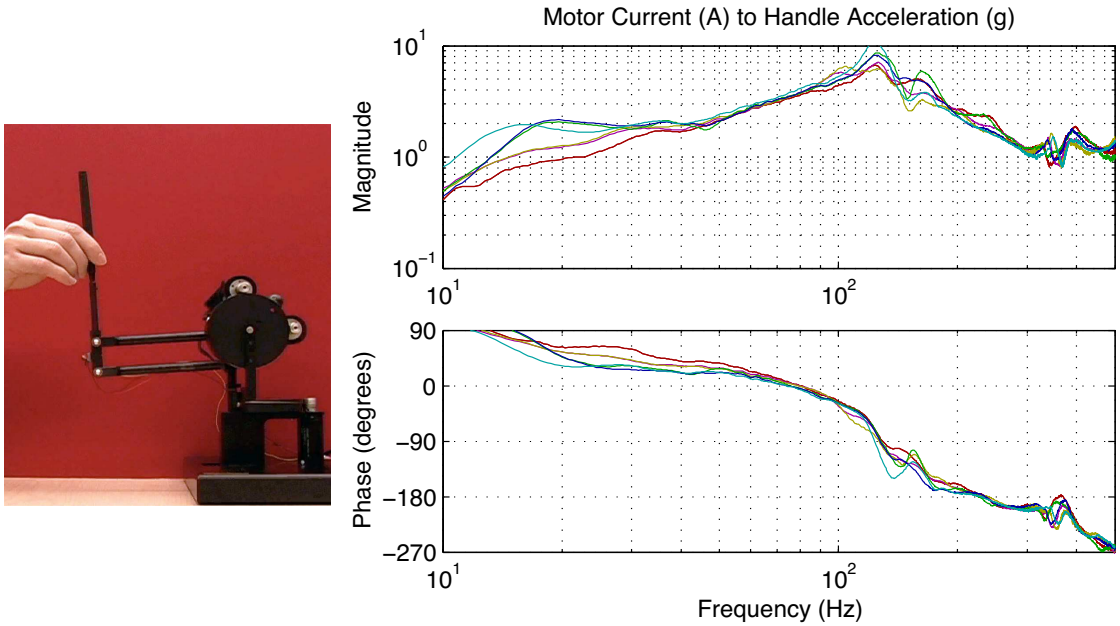

Figure 3.8: The Phantom and its ETFEs for handle acceleration in the vertical direction.

3.4.3 Determining Linearity and Time Invariance

Performing the above procedure at a variety of signal magnitudes, frequency ranges, positions in the haptic interface's workspace, and user grip force levels will elucidate the linearity of the system's behavior. If the experimental Bode plots match closely for a variety of conditions, the haptic interface's mechanical dynamics can be modeled as independent from such changes. Ideally, all of the comprehensive ETFEs would match, and a linear model could be tuned to match the total observed behavior.

One good example of a system that behaves in a fairly linear, time-invariant manner is the vertical handle acceleration output of the Phantom shown in the left half of Figure 3.8. This system was identified via comprehensive evaluation. A user passively held the stylus with a moderate grip force while a 2.5-second-long swept-sinusoid current command from ten to 500 hertz was applied to the system. An accelerometer mounted near the user's fingers enabled the recording of handle acceleration. Four tests were performed at three current command magnitudes for two users, and results were averaged across tests. The frequency content of input and output signals was compared by taking the ratio of their discrete Fourier transforms (DFTs) in the complex domain, as detailed above. The resulting curves are shown in the right half of Figure 3.8. The striking agreement across input

signal magnitude and frequency order indicates that this haptic interface system behaves in a manner that is approximately linear and time invariant; some deviations are observed below 30 hertz, where user intention and differences between users affect the behavior of the system more strongly. Overall, though, the dynamics of this haptic interface appear consistent, boding well for our ability to model them with an LTI system.

For some haptic interfaces, the comprehensive ETFEs generated under various test conditions do not match well and/or do not appear to correspond to an easily deduced linear system. These discrepancies can stem from higher-order dynamics, nonlinearities, and the influence of the user, which can all be difficult to characterize from the single-input single-output frequency-domain perspective of comprehensive evaluation. In these cases, building a model step by step from the physical system, as explained in Section 3.5, can help one thoroughly understand its behavior.

3.4.4 Fitting an LTI Model

A linear, time-invariant model can be fit to a set of ETFE results by hand-tuning the placement of poles and zeros; the model developed for vertical handle acceleration on the Phantom is shown as the black dashed line in Figure 3.9. At steady state, one would expect the user-master system to behave like a spring, providing the transfer function from requested current to handle acceleration with two zeros at the origin. Although identification of low-frequency dynamics is obscured by the user's reaction to the input, one does observe an approximate slope of plus two at ten hertz, along with a phase lead of about 90 degrees. For this model, which is concerned primarily with behavior above 20 hertz, a single zero at five hertz was deemed sufficient, followed by a pole at 25 hertz; as will be discussed in Chapter 4, models to be used in high-frequency acceleration matching should not contain zeros at the origin for $\hat{G}_m(s)$ because they become pure integrators when the transfer function is inverted.

The system's primary resonance occurs at 129 hertz and is followed by alternating pairs of lightly damped zeros and poles. The full model is a seventh-order, relative-degree-four transfer function with a 0.25 millisecond time delay. When fitting a linear model to ETFE results, both magnitude and phase should be considered, adding in a small time delay to further depress the phase at high frequency. If magnitude and phase cannot both be matched, the system has underlying nonlinearities and should be identified using successive isolation. If the ETFE curves from different testing configurations diverge slightly, the

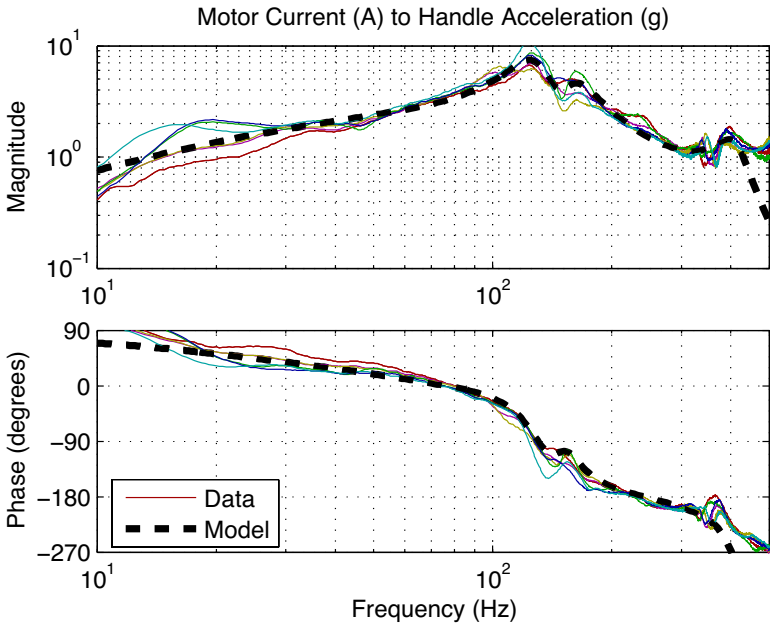

Figure 3.9: Hand-fit linear time-invariant model for the system shown in Figure 3.8.

future use of the model should guide the tuning. If it is to be employed in high-frequency acceleration matching, I fit the model to match the curves with the highest magnitude to avoid over-stimulating resonant behavior when the model is inverted. If it is to be used in cancellation of induced master motion, I tune the model to match the lower magnitude curves to avoid over-canceling. If a general understanding of the device's behavior is desired, the model can be matched to the average of the curves, or successive isolation can be employed to determine the underlying cause of the variations.

3.4.5 Time-Domain Validation

Once a model has been formulated, its time-domain response must be validated for a set of typical input transients. For example, the system model obtained in the previous section was validated in the time domain by playing a variety of transients as users held the stylus with a moderate grip force and tapped on a virtual object. Three sample current and acceleration traces are shown in Figure 3.10, and one sees that the model's response matches the measured values very closely, especially for the short durations characteristic of contact transients. I conclude that a simple, user-invariant dynamic model is useful in describing the vertical acceleration response of the Phantom during haptic interactions, provided the user

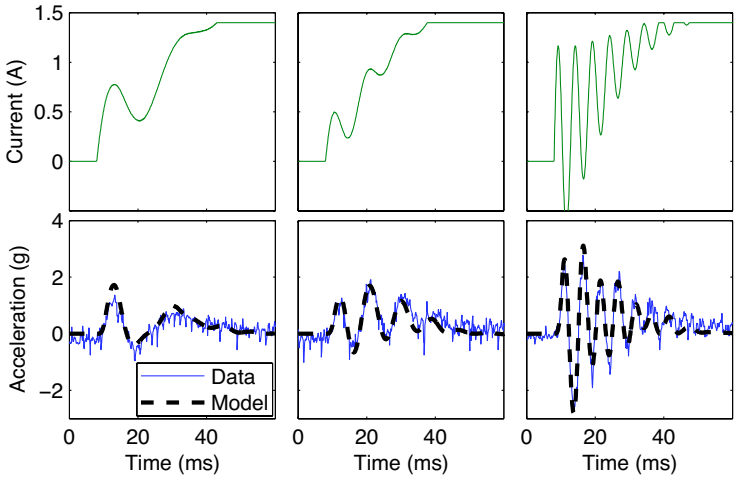

Figure 3.10: Time-domain verification of the LTI model for Phantom handle acceleration.

maintains a consistent grip on the stylus. The influence of grip force on system dynamics is discussed in the next section. Though it is sometimes difficult to interpret their individual parameters, models obtained through comprehensive evaluation can aptly and efficiently capture a system's frequency response under a range of conditions.

3.4.6 Representing the Influence of Grip Force

Although many haptic interfaces can be treated as LTI systems, they always contain an element that varies over time: the user. Changes in hand configuration and finger positioning can significantly change the effective impedance of the user's hand at the endpoint of the device. And when a user changes the force with which he or she grasps the handle, the stiffness and damping of his or her arm muscles increase [84]. These changes affect the transmission dynamics of the entire system and can alter the response significantly, especially near resonant frequencies. Although dynamics can be assumed constant if the user maintains a consistent grip, these variations must generally be taken into account for best performance.

Comprehensive evaluation can be extended to create an accurate model of a haptic system under varying grip force. A series of tests is run at different grip force levels, using a small force sensor on the handle to measure the user's present grip force. I conducted this type of evaluation on the one-degree-of-freedom interface shown in Figure 3.11, aiming

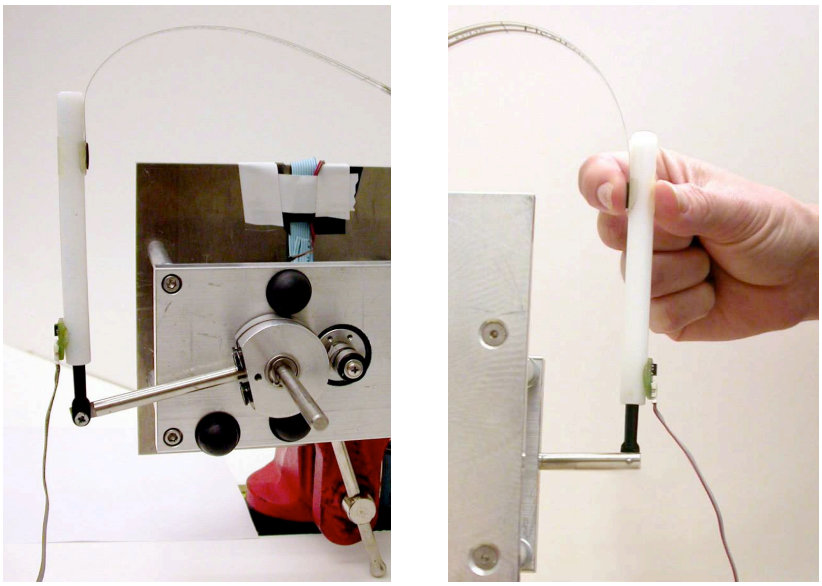

Figure 3.11: Two views of the custom one-dof haptic interface.

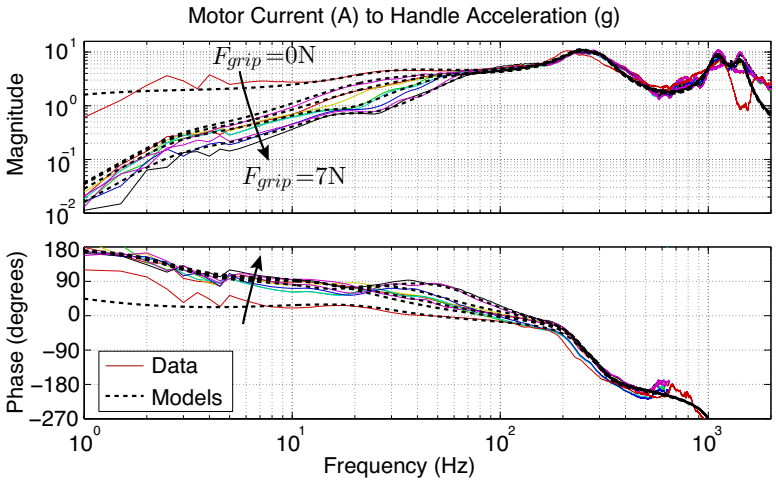

Figure 3.12: The one-dof system's grip-dependent ETFEs with representative dynamic models.

to characterize the grip-dependent dynamics between its current command and stylus acceleration. This interface uses an RE025 Maxon motor taken from a Phantom, and it has approximately the same amplification factor from motor torque to stylus force as the Phantom interface previously characterized. This interface was studied to determine whether comprehensive evaluation techniques would work well on non-standard haptic interfaces. A high-bandwidth accelerometer is attached to the lower portion of the stylus to sense vertical acceleration, and a force-sensing resistor from FlexiForce [40] is located beneath the user's fingers for sensing grip force.

Each solid curve shown in Figure 3.12 is the result of a complex-domain average of four independent, two-second long, swept-sinusoid trials from one to 2000 hertz. The amplitude of the input sinusoid was fixed at 0.75 amps, with the starting polarity reversed for two of the four trials. This amplitude was chosen because it is in the middle of the motor's steady-state current range, producing approximately 0.725 newtons at the user's hand when applied continuously. Grip force for each curve was maintained near a constant value throughout each test using a real-time graphical display for user feedback; trials over which grip force varied more than ± 0.5 newtons were discarded and repeated. The zero-Newton grip force case represents the natural dynamics of the device, which must also be characterized.

The five dashed-line models shown in Figure 3.12 were hand-tuned to closely match the ETFE results without excessive dynamic complexity. A basic approximation of the dynamics below 100 hertz shows that when the user is disconnected from the stylus (i.e., zero grip force), the low frequency dynamics are dominated by the inertia of the stylus and mechanical linkage. When the user holds on, one observes the addition of a low-frequency spring and a mid-frequency damper. These additional elements cause the frequency response to have an approximate slope of +40 dB/decade below five hertz and +20 dB/decade between five hertz and 80 hertz. It appears that changes in the stiffness of this spring and dissipation in the damper are the major effects of changing grip force, where a decrease in ETFE magnitude corresponds to an increase in the respective parameter value. Many researchers have identified the trend of increasing damping and stiffness with increasing grip force [15, 58, 147], which is generally supported by the ETFE results. Once LTI models have been obtained for a set of grip force levels, the dynamics of the device at intermediate values can be approximated via interpolation between neighboring models.

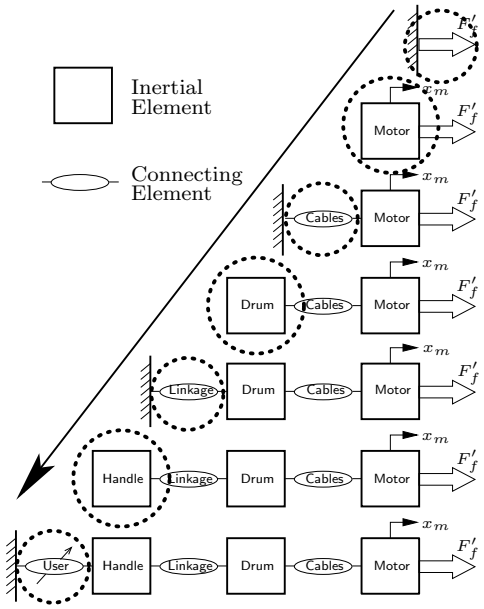

Figure 3.13: Successive isolation of the user-master system.

3.5 Successive Isolation

While comprehensive evaluation yields a good model for many systems and can even be extended to account for variations in user grip force, it is sometimes necessary or preferable to obtain a model that has physically based parameters rather than poles, zeros, and a gain; such situations can be handled via successive isolation. This approach is especially applicable if the user-master system exhibits strong nonlinearities or time-varying behaviors, as it helps one identify an appropriate model structure. Furthermore, successive isolation can be useful on systems that contain many electronic and mechanical elements; attempting to identify such a complex system all at once makes it hard to distinguish overlaid effects, each of which may be simple and based on a well-understood phenomenon. A system like this can be characterized most effectively by breaking its dynamic chain at points progressively farther away from the motor's current input and position output, as illustrated in Figure 3.13. For each mechanical configuration, different current signals such as sine waves or step inputs are commanded to the system, and the resulting motor motion is recorded. Such a strategy allows successive isolation of each element and its individual contribution to system behavior. The procedure for successive isolation is described below and then demonstrated on a test system.

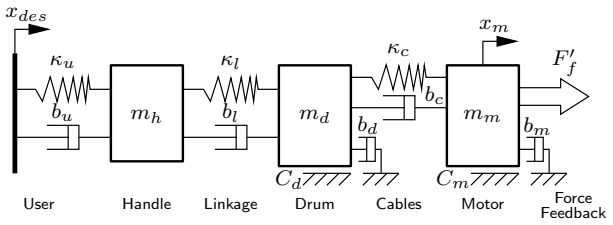

Figure 3.14: Sample parameterization for a sixth-order master model.

3.5.1 Procedural Overview

An accurate system model can be obtained by identifying each of the haptic interface's many dynamic elements in turn, as originally proposed in [81]. Such an approach yields a model whose parameters have physical significance, facilitating the addition of a varying user element and informing the processes of device design and modification. I begin by separating the system's dynamic chain into alternating inertial and connecting elements, as illustrated in the bottom row of Figure 3.13. This model structure is shown with sample constitutive elements in Figure 3.14, using κ to stand for possibly nonlinear stiffness, b for linear viscosity, C for Coulomb friction (discussed below), and m for mass. Once the system is divided into inertial and connecting elements, I proceed to identify each element's properties in turn, mechanically breaking the dynamic chain at points progressively farther away from the motor's force command input and position output, which are F_f and x_m respectively. At each stage, I conduct quasi-static and dynamic tests, comparing observed and modeled behavior in both the frequency and time domains; only the model's outermost parameters are adjusted to make its response match the present round of observations.

The first dynamic element to consider is the current amplifier, which is commonly assumed to be ideal. The controller's current command is often treated as a command of effective steady-state force at the user's hand, as the two are related by the proportionality

$$F_f = \frac{\rho\, k_t}{h}\, i_{cmd}, \tag{3.17}$$

where h is the distance from the drum axis to the center of the hand and ρ is again the unitless gear ration defined in (3.4). While F_f is the commanded force, I define F'_f to be the force actually applied to the system by the amplifier's true current output, which may incorporate a variety of dynamic relationships such as a low-pass filter and saturation. Next successive isolation starts the process of adding inertial and connecting elements in order,

working away from the amplifier, as illustrated in Figure 3.13. When an inertial element is appended to the dynamic chain, the connecting element beyond it must be disengaged so that the mass can move freely; for example, when the motor is added, the cables must be disconnected. I then seek to determine the new element's inertia and dissipation parameters.

Although dynamic models typically include only viscous friction, a Coulomb friction model is more appropriate for haptic interfaces due to their relatively low operating velocities [32]. Coulomb friction is defined as

$$C = \begin{cases} -c & \text{if } \dot{x} > 0 \\ 0 & \text{if } \dot{x} = 0 \\ +c & \text{if } \dot{x} < 0, \end{cases} \tag{3.18}$$

where $\dot{x}$ is the relative interface velocity of the body that receives the force C. Many more sophisticated methods for modeling static and dynamic friction have been developed that could also be incorporated into successive isolation, e.g. [60], but this simple model works well for characterizing the behavior of most haptic interfaces. Nonlinear friction like that described by the Coulomb model is commonly found at sliding interfaces, and its effect is characterized by linear, rather than exponential, decay of transient dynamics. Assuming (3.18) approximately holds for a given system, the magnitude of Coulomb friction can be determined by measuring the force required to move the present set of connected inertial elements across the workspace under closed-loop position control.

The mass and the viscous friction coefficient of an inertial element can be highlighted by adding a virtual spring to the motor's end of the chain and analyzing both the frequency and time response of the resulting system. When identifying inertial elements, open-loop experiments that do not include a virtual grounding force should generally be avoided; the unconstrained movements that result are very sensitive to initial conditions and are not representative of the high-frequency oscillatory behavior that I wish to understand. The focus advances to the next connecting element once an inertial element's parameters are estimated and validated.

The connector should then be attached to its neighboring masses, and the distant mass should be immobilized mechanically; for example, when the cables are first added to the system, the drum's position should be locked. The important characteristics for each connecting element are its possibly nonlinear stiffness and its dissipation. Although dynamic models commonly include only linear springs, large deflections of most physical objects

reveal nonlinear relationships between position and force as well as occasional hysteresis. The stiffness of a connecting element can be examined by slowly varying the motor force and measuring the resulting position variations. Non-constant slope can be appropriately parametrized, and hysteresis can be modeled with a new method I developed, which is presented below in Section 3.5.3. A connecting element's properties can also be highlighted using the ETFE method presented above, and its overall behavior can be validated via time response tests. When a connection is fully characterized, the inertial element beyond it is added to the chain and the successive isolation process continues.

3.5.2 Sample Implementation

Through application to a one-dof testbed, this section demonstrates the methodology of successive isolation for characterizing the behavior of a haptic device. After a comprehensive evaluation of its dynamics under various conditions, the user-joystick system pictured in Figure 3.15 was found to exhibit a higher-order, nonlinear response that required successive isolation. A detailed model of its dynamics is developed through careful application of a variety of system identification techniques. The model is constructed around translation of the user's hand, converting torques and rotations to their effective values in this space.

The hardware of this system is an Impulse Engine 2000, a research-quality force-feedback joystick produced by Immersion Corporation. As shown in Figure 3.15, each of its two axes includes a DC motor connected to the handle via a cable drive and a mechanical linkage; the base of the joystick contains the motors and their respective current amplifiers. This

Figure 3.15: Immersion Impulse Engine 2000 joystick.

investigation centers on the forward/backward degree of freedom, keeping the left/right axis in the center of its travel. The joystick is controlled via an ISA card by a personal computer running RTAI Linux, a hard-real-time open-source operating system [127]. All device-control software was custom written in C, taking the place of the software development kit provided by the manufacturer.

Each element of this master system, from the timing of the servo loop to the effective impedance of the user, was tested in succession, using Matlab and Simulink for data analysis and model simulation. An increasingly complex system model was constructed during this isolation process, changing only the model's outermost parameters to match simulated to experimental behavior at each step. The objective of the simulation is to provide an accurate discrete-time model of the haptic interface's relevant behavior using physically based and individually determined parameters. At the conclusion of this successive isolation demonstration, Section 3.5.4 provides the full identified model with all parameter values.

Servo Loop Timing

Digital controllers are designed to execute at a steady frequency, measuring the device's present position and computing the appropriate feedback force at each servo cycle. While important for accurate velocity estimation during normal use, predictable servo timing is crucial for accurate system identification and model-based control. Servo loop consistency can be observed by recording the actual time at which the system's control code starts and finishes for each cycle. This measurement can be performed using a low-level timer internal to the computer such as the processor's time stamp clock. It can also be done by setting and clearing a digital output line such as a parallel port bit and watching the output with an oscilloscope. Control code should execute at its specified rate without missed cycles, drift, or variations in period. If these behaviors are unavoidable due to computer architecture, the controller must record the actual time elapsed between cycles using a low-level timer, and these measurements must be used during system identification and model-based control. The RTAI Linux platform of the joystick testbed achieves uniform servo loop timing at frequencies up to and beyond ten kilohertz, permitting a fixed rate assumption.

Current Amplifier

Although usually assumed to provide perfect signal tracking, the current amplifier can add slight nonlinearities and important model dynamics to the haptic system. First, its

steady-state gain and zero offset should be calibrated in hardware and/or software to ensure that DC current commands from the computer are accurately applied to the motor, which facilitates the identification process. The maximum steady-state current supported by the motor should also be enforced and included in the model. The gain and offset of the Impulse Engine amplifier were calibrated using a high-precision ammeter and corrected in software.

Second, the frequency response of the amplifier should be characterized while the motor position is mechanically locked. Sinusoidal current commands up to one kilohertz should be tested, keeping the servo rate at least a factor of ten above the test frequency. The voltage across a sense resistor in series with the motor can easily be monitored with an oscilloscope to detect high-frequency attenuation. More complete analysis of both magnitude and phase can be achieved by measuring current in real time with a sufficiently fast analog-to-digital computer card. With such computer monitoring, a swept sine wave with linearly varying frequency can be commanded instead of signals at discrete frequencies. The system's behavior can then be examined through an ETFE. First- or second-order low-pass filtering will often be elucidated through this investigation; observed behavior should be characterized and included in the system model, for it will strongly affect the device's response to high-frequency force feedback commands. The joystick testbed's linear amplifier responds excellently past one kilohertz, showing only the fixed delay of the servo cycle at all frequencies.

A third phenomenon to note is that of back-EMF; rotation of a DC motor creates an opposing voltage that the amplifier must attempt to overcome. The fixed voltage of its supply limits the current an amplifier can source when the motor is rotating quickly. High motor velocities are seldom encountered during normal operation, but they may occur when identifying mechanical resonances, thereby attenuating the applied current. This effect can be modeled using knowledge of motor and amplifier parameters, or it can be circumvented by measuring current flow during all identification tests. The joystick testbed displays this type of current attenuation only at abnormally high motor velocities with commands near maximum current, so I opted to measure current rather than build a model of this peripheral effect.

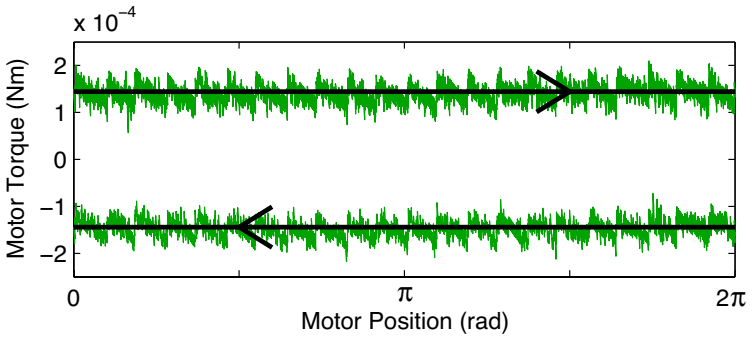

Figure 3.16: Torque required to move the isolated motor forward and backward under position control.

Motor and Encoder

The motor shaft, the first moving element in the master's dynamic chain, can be isolated by disconnecting the cables that couple it to the drum, leaving it free to move. The fundamental parameters of motor torque constant and encoder resolution may be measured or taken from appropriate manufacturer data sheets. The actual torque constant and the relative spacing of encoder ticks usually match these given values very closely. The Coulomb and viscous friction parameters of the motor, c_m and b_m, can then be identified using a closed-loop position controller. Recording the torque required to slowly rotate the motor in each direction provides a measure of its frictional losses. Performing this test at a range of speeds differentiates Coulomb friction from velocity-dependent viscous friction. Unlike gradually increasing the torque until the motor moves, this test also illuminates any position-dependent friction such as that which results from damaged motor bearings. The testbed's Maxon RE025-055-35 motor exhibits Coulomb friction only, varying minutely with commutator-brush position as illustrated in Figure 3.16. The twenty-two sloped sections are caused by its two brushes transitioning between the commutator's eleven segments, a variation that can be disregarded as insignificant.

The motor shaft's total inertia can be estimated by summing the rotational inertias given on motor, encoder, and capstan data sheets. The motor's effective mass in handle space, m_m, can be estimated from these values as

$$m_m = (J_m + J_{enc} + J_{cap}) \cdot \rho^2 \cdot \left(\frac{1}{h}\right)^2, \tag{3.19}$$

where J signifies rotational inertia, ρ is the unitless gear ratio between motor and drum rotation, and h is the distance from the drum axis to the chosen tip location on the handle. This value can be confirmed by conducting a set of motor step responses under closed-loop position control, which acts as a virtual spring to ground and highlights the inertia via lightly damped oscillations. Starting with a model of the discrete servo cycle, the motor's identified friction, and the position controller, I fine-tuned the simulation's motor to match the system's observed time response and achieved excellent agreement.

Cables

The dynamic properties of the cables can be isolated by re-connecting them to the motor and locking the position of the drum. The cable should be tight enough to prevent relative motion between these two elements, but over-tightening increases the friction at this interface. The possibly nonlinear stiffness, κ_c, can first be investigated by slowly varying

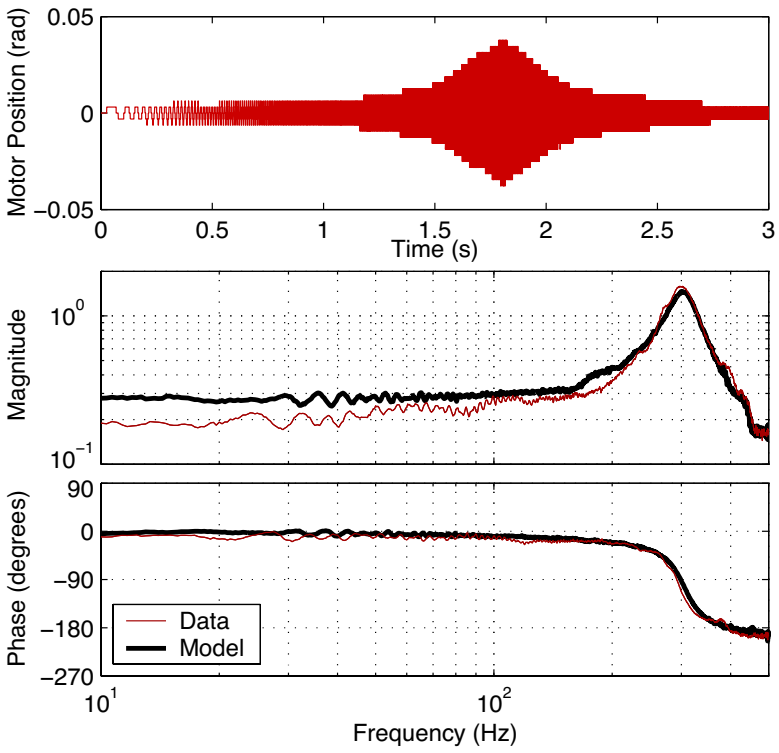

Figure 3.17: Resonant response of the cable/motor assembly.

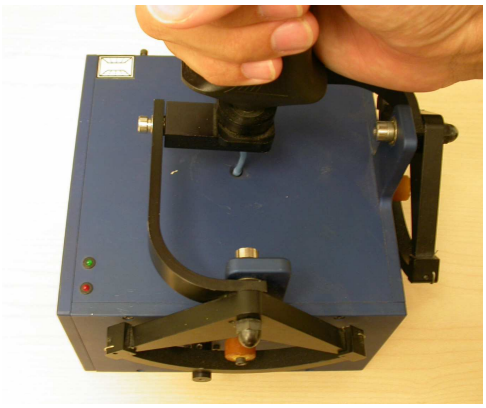

Figure 3.18: Top view of the joystick testbed.

the motor torque across its full range using a triangle or sine wave, recording the resulting position waveform. The cables will generally allow only a few encoder ticks of movement in either direction, providing only a rough stiffness estimate, but the resulting torque/position plot may show an enclosed area, the hallmark of hysteresis.

To augment this measurement on the joystick, I applied a swept-sine-wave input that varied from ten to 500 hertz to both the testbed and the simulation. Figure 3.17 shows the actual motor's time response, clearly demonstrating a resonance, as well as an empirical transfer function estimate (ETFE) of both the data and the simulation after parameter fitting. Although the system's Coulomb friction and hysteresis make it nonlinear, this technique can be useful in understanding and modeling its resonant behavior. A small amount of hysteresis was observed in the cables of the test system, which changes their frequency response slightly based on input magnitude, but this effect was deemed negligible and appropriate values were chosen for the linear cable stiffness, k_c, and linear cable damping, b_c. The effect of hysteresis on dynamic response is discussed below in detail, regarding the joystick's linkage element.

Drum

The drum's inertial and frictional characteristics can be identified in much the same way as those of the motor. Disconnecting the linkage leaves the drum free to spin; on the joystick testbed, the linkage was separated where the bent member joins the beam from the base of the handle, splitting its small mass between the drum and the handle. This disconnection point is visible in Figure 3.18 as the silver screw head at the upper left corner of the

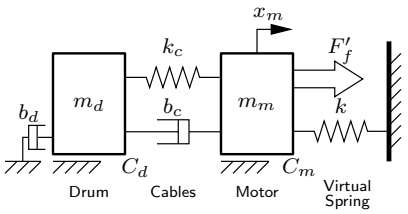

Figure 3.19: Testing configuration for drum parameter identification.

linkage. As with the motor, open-loop experiments that did not include a virtual grounding force were avoided. Instead, the drum's parameters were identified by adding a position controller to the motor, grounding that end of the dynamic chain via a virtual spring as illustrated in Figure 3.19. Step responses with various magnitudes and with several controller position gains were then performed via F_f to elucidate the system's nonlinear behavior. A combination of hand-fitting and nonlinear unconstrained optimization was used to fit the parameters of drum mass (m_d), viscous friction between drum and ground (b_d), and Coulomb friction between drum and ground (c_d); an initial set of parameters was estimated by hand, and the Matlab function `fminsearch` was employed to find a locally optimal parameter set, using a sum of squared position errors to quantify the accuracy of particular simulations on each test. These parameters could also have been identified by moving the system across its workspace under position control at various speeds, but the small workspace of the drum (about 120°) made such tests difficult.

Linkage

After locking the handle position and ensuring that there is no backlash in any of the system's connections, the linkage can be examined using a force/position plot and a frequency response. The testbed's linkage was found to be an order of magnitude softer than the cables, and it displayed a rate-independent hysteresis, indicated by a significant enclosed area on its force/position plot. This hysteresis caused the system's resonant frequency and magnitude to depend strongly on input magnitude, a behavior that cannot be captured with a linear or position-dependent stiffness model. Unwilling to disregard this effect, I developed a new method for modeling hysteresis in dynamic systems, which is described in detail in Section 3.5.3.

The new hysteresis model was added to the built-up simulation to represent the linkage stiffness. The testbed's feedback force was slowly varied, and the motor's resulting position

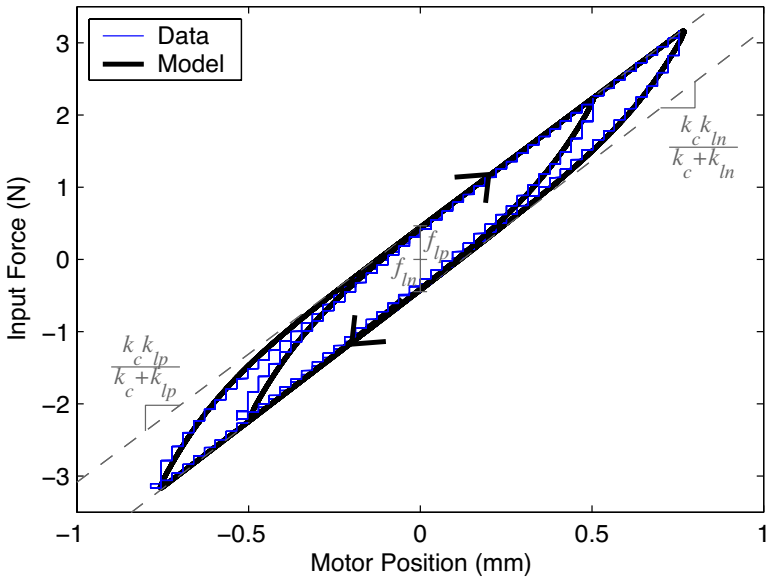

Figure 3.20: Hysteretic behavior of the cables and linkage in series.

change was recorded, as shown in the force/position plot of Figure 3.20. Though the slope is approximately constant, significant hysteresis is observed via the enclosed area of the curve, which takes the same shape for a wide range of forcing frequencies. I deduced the appropriate parameters for my new hysteresis model from the experimental data's positive and negative velocity asymptotes, which represent the series stiffness of linkage and cables, and the shape that the force/position trace takes when transitioning between asymptotes. See Section 3.5.3 for a full explanation of the five linkage parameters, k_{lp}, k_{ln}, f_{lp}, f_{ln}, and γ. After parameter fitting, this simple model of rate-independent hysteresis aptly characterized the system's observed behavior.

A good hysteresis model is important because apparent stiffness strongly depends on deflection magnitude. Small deflections give a relatively stiffer appearance, so the system's resonant frequency changes with signal amplitude. Swept-sine-wave inputs of different magnitudes were applied to both the actual and simulated testbeds, with a sample result shown in Figure 3.21. These tests enabled the fitting of the linkage's viscous friction parameter, b_l. After this choice, the observed and simulated ETFEs matched well, responding similarly for a range of input magnitudes. This linkage model enabled me to accurately characterize the haptic interface's response to high-frequency force feedback, as the shifting resonance cannot be adequately approximated with a linear model.

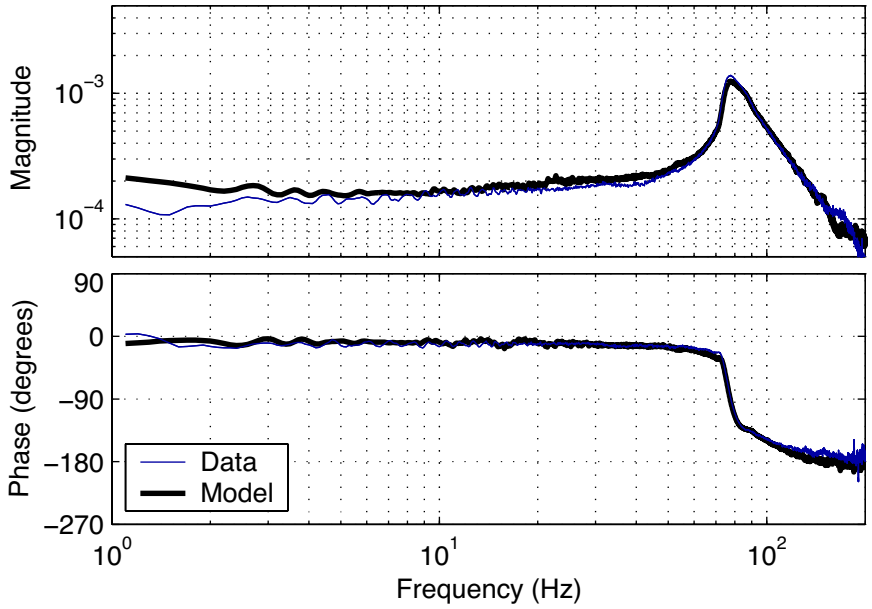

Figure 3.21: Nonlinear resonant response of the assembled motor, cables, drum, and linkage.

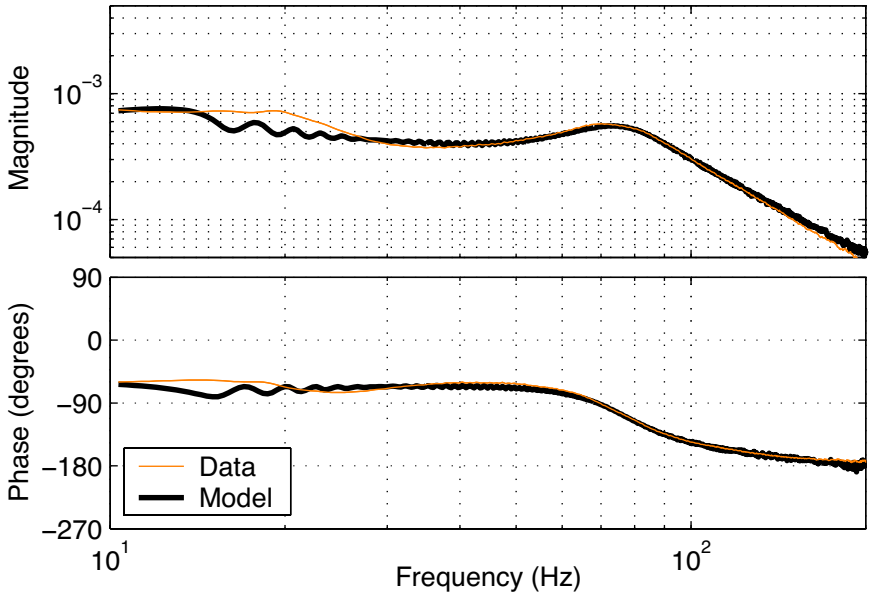

Figure 3.22: Response of the complete system when held by a user.

Handle and User Impedance

The last elements of the master's dynamic chain, the mass of the handle and the impedance of the user, can be characterized in the frequency domain. Swept sine waves varying from 10 to 200 hertz were applied while a user held the joystick in a comfortable grasp. Figure 3.22 shows the corresponding ETFE, to which the full model was matched via selection of handle mass, m_h, user stiffness, k_u, and user damping, b_u. The subdued resonance at 70 hertz stems from the handle and mechanism vibrating against the flesh of the user's hand. It was found not to change significantly with variations in grip force, which can likely be traced to the long distance between the handle and the motor.

If this model needed to accurately predict handle acceleration as well as motor movement, accelerometer readings could be recorded during this round of testing and an appropriate relationship between grip force and user stiffness and damping could be determined. A study that I conducted in 2003 showed approximately linear variation of the rotational stiffness and damping of the user's wrist joint in flexion and extension under changes in grip force [84]. At a minimum, these parameters can be made to take the value of zero when the user is not holding the device and to increase linearly from there, though this approach does necessitate a grip force sensor. If detecting the presence of the user was important, another option would be to use a momentary contact switch as a binary grip force sensor so that the system can select between the non-user and user models.

3.5.3 Modeling Hysteretic Stiffness

A rate-independent hysteresis such as the one identified in the joystick's linkage can be attributed to combined elastic and plastic deformation, which is generally modeled using a play/stop operator [93]. Unfortunately such a model cannot be used in dynamic simulation because it maps force to position rather than position to force. A position-to-force model is necessary for forward simulation of haptic interface dynamics because the hysteretic element must be able to accept as input the position difference between two masses and output the resulting force between them. This is the role played by a standard spring – serving as a feedback element around the double integrators of the two connected masses. Furthermore, the use of a play-stop operator produces an output that does not vary smoothly. In order to enable dynamic simulation of compliant elements with hysteresis, I developed a new model structure that takes an input deflection and produces a smooth, hysteretic force signal.

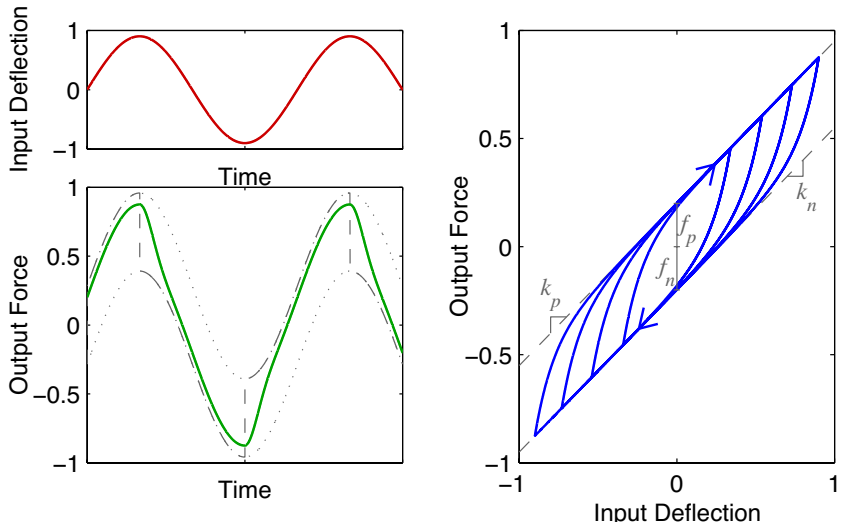

Figure 3.23: New hysteresis model relating input deflection to output force.

As seen in Figure 3.20's plot of linkage behavior, cycling the force applied across a stiffness element with hysteresis traces a clockwise path in the force/position plane, enclosing an oblong shape with pointed upper-right and lower-left corners. A non-hysteretic spring would describe a single diagonal line rather than enclosing an area. Hysteretic behavior is illustrated for a generic case in Figure 3.23, showing the single time-domain plots of deflection and force, as well as these two signals plotted against each other for several input magnitudes. In this plot, deflection is treated as the input, as that is the type of model we need for dynamic simulation, but the behavior is the same as that exhibited by the linkage. The observed relationship between position and force can be described by the two limits

$$\lim_{x \to +\infty} F(x) = f_p + k_p x \tag{3.20}$$

and

$$\lim_{x \to -\infty} F(x) = f_n + k_n x, \tag{3.21}$$

where F is the force across the element and x is its deflection from nominal. I have defined f_p and f_n to be the vertical offsets of the positive and negative asymptotes, respectively, and k_p and k_n to be their slopes. The f parameters have units of force, and the k parameters have units of force per distance like standard stiffnesses; the parameters f_p, k_p, and k_n are positive valued, and f_n is negative. The asymptotic relationships described by (3.20) and (3.21)

mathematically describe the target behavior for our model, along with the requirement for smooth transitions between asymptotes, and they suggest a way to construct a model for discrete real-time simulation.

The developed model produces the desired output by switching between two offset stiffness lines based on positive or negative position changes and smoothing the output with a tunable spatial filter. The coarse, pre-filtered force estimate is computed by

$$\hat{F}_i = \begin{cases} f_p + k_p x_i + g_p & \text{if } x_i > x_{i-1} \\ \hat{F}_{i-1} & \text{if } x_i = x_{i-1} \\ f_n + k_n x_i + g_n & \text{if } x_i < x_{i-1} \end{cases}, \tag{3.22}$$

where i is the index of the current simulation cycle and g_p and g_n are supplementary force offsets that are necessary for convergence to the desired limits, as discussed and defined below. This equation produces a coarse force estimate, $\hat{F}$, that always lies on one of the two asymptote lines and switches abruptly when the element's deflection changes from increasing to decreasing or the reverse. If the deflection has not changed since the previous simulation cycle, the element's force output remains the same.

The rough force estimate provided by (3.22) must be smoothed to match experimental observations. As the hysteresis is independent of the rate at which it is excited, the smoothing cannot be done in the time domain. Instead, I limit changes in force based on the deflection that has transpired since the last simulation cycle. A first-order low-pass filter serves this purpose beautifully. In continuous space, this low-pass filter would take the form

$$\frac{\hat{F}(s)}{F(s)} = \frac{\gamma}{s+\gamma}. \tag{3.23}$$

The parameter γ is the spatial filter bandwidth; it has units of radians per distance, as does the Laplace operator s in this case. When implemented in a discrete simulation, this filter can be represented by

$$F_i = \frac{\gamma X}{1+\gamma X}\hat{F}_i + \frac{1}{1+\gamma X}F_{i-1}, \tag{3.24}$$

with the change in deflection X defined as

$$X = |x_i - x_{i-1}| \tag{3.25}$$

rather than the usual T employed in temporal filters. Such a choice produces a consistent

shape in the force/position plane regardless of the time derivative of deflection, and the smoothness can be tuned by adjusting γ. Smaller values of gamma produce smoother transitions between the two curves, as they limit force changes with respect to deflection more severely.

Implementing the above model without carefully choosing g_p and g_n does not yield the desired force/position behavior: rather than converging to the limits described by (3.20) and (3.21), the system asymptotically approaches lines that are each offset towards the origin of the plot. This discrepancy can be traced to the steady-state error of the first-order filter. The Laplace transform of its error, $E(s)$, can be formulated in continuous space as

$$E(s) = \hat{F}(s) - F(s) = \left(1 - \frac{\gamma}{s+\gamma}\right)\hat{F}(s) = \frac{s}{s+\gamma}\hat{F}(s). \tag{3.26}$$

When the changes in deflection flip sign and the raw force signal jumps to the other asymptote, the filter sees a combined step and ramp input. The position-based Laplace transform of this raw-force input for a switch from lower to upper asymptote is

$$\hat{F}(s) = \frac{f_p - f_n}{s} + \frac{k_p}{s^2}. \tag{3.27}$$

The first term in (3.27) represents the vertical step from lower to upper asymptote, and the second term represents the sloped progress thereafter.

The steady-state error to this input can be determined using the final value theorem:

$$\lim_{x\to+\infty} E(x) = \lim_{s\to 0} s\frac{E(s)}{\hat{F}(s)}\hat{F}(s). \tag{3.28}$$

The transfer function from coarse force estimate to filter error can be determined from (3.26) and substituted into (3.28) along with the Laplace transform of the input, providing

$$\lim_{s\to 0}\left[(s)\left(\frac{s}{s+\gamma}\right)\left(\frac{f_p - f_n}{s} + \frac{k_p}{s^2}\right)\right] = \lim_{s\to 0}\left[\frac{s(f_p - f_n)}{s+\gamma} + \frac{k_p}{s+\gamma}\right] = 0 + \frac{k_p}{\gamma} \tag{3.29}$$

The filter's force output is thus seen to exhibit a steady-state error of k_p/γ when tracking the positive asymptote. It can also be shown that the filter has a steady-state error of $-k_n/\gamma$ when approaching the negative asymptote, though the derivation is omitted here for brevity. I correct for this effect by offsetting the asymptotes by the error amount via the g

terms in (3.20) and (3.21), as follows:

$$g_p = k_p/\gamma \tag{3.30}$$

and

$$g_n = -k_n/\gamma. \tag{3.31}$$

As the k terms have units of force per distance and γ is in radians per distance, g_p and g_n have units of force, as expected. The net effect of this compensation is to move the raw force's positive asymptote up and its negative asymptote down to ensure convergence to the desired lines after spatial filtering. With the corrective g terms in place, the filtered force gradually converges to $f_p + k_p x$ when the derivative of deflection is positive, and it converges to $f_n + k_n x$ when the deflection derivative is negative. To my knowledge, such a hysteresis model has not before been developed, and it can be applied to compliant elements found in many dynamic systems.

3.5.4 Full Model

A step-by-step account of the joystick identification process is provided above, and this section presents an overview of the final model. The servo-loop timing was found to achieve sub-micro-second accuracy, and the haptic interface's linear current amplifier behaved almost ideally, giving $F'_f = F_f$ for full-scale current sinusoids at frequencies up to 1000 hertz. The mechanical elements of the system are well-approximated by the nonlinear lumped-parameter model shown in Figure 3.24.

Motor, drum, and handle movements are mapped to translation of the user's hand by the unitless motor-to-drum gear ratio, $\rho = 15$, and the distance from the drum axis to the handle endpoint, $h = 0.15$ m. Both the motor and the drum exhibit significant Coulomb friction, and the cables behave like a simple spring and damper. Modeling the spring force of the joystick's linkage, F_{κ_l}, requires the nonlinear hysteretic description developed Section 3.5.3. All of the model's parameters are provided in Table 3.1, and its basic equations of motion follow.

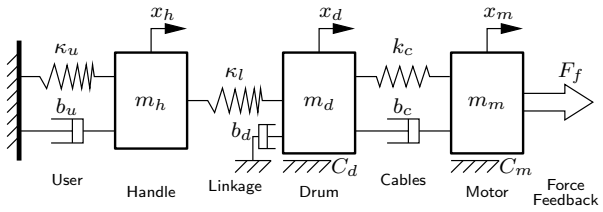

Figure 3.24: Nonlinear lumped-parameter model of the user-master system using κ symbols to signify nonlinear stiffness relationships.

Table 3.1: Identified joystick model parameters.

Parameter	Value	Parameter	Value
k_u	40 N/m	m_d	0.0075 kg
b_u	23 Ns/m	b_d	0.02 Ns/m
m_h	0.014 kg	c_d	0.01 N
k_{lp}	3900 N/m	k_c	39500 N/m
k_{ln}	4000 N/m	b_c	3.6 Ns/m
f_{lp}	0.465 N	m_m	0.01075 kg
f_{ln}	-0.445 N	b_m	0 Ns/m
γ_l	8500 rad/m	c_m	0.0144 N
b_l	0 Ns/m		

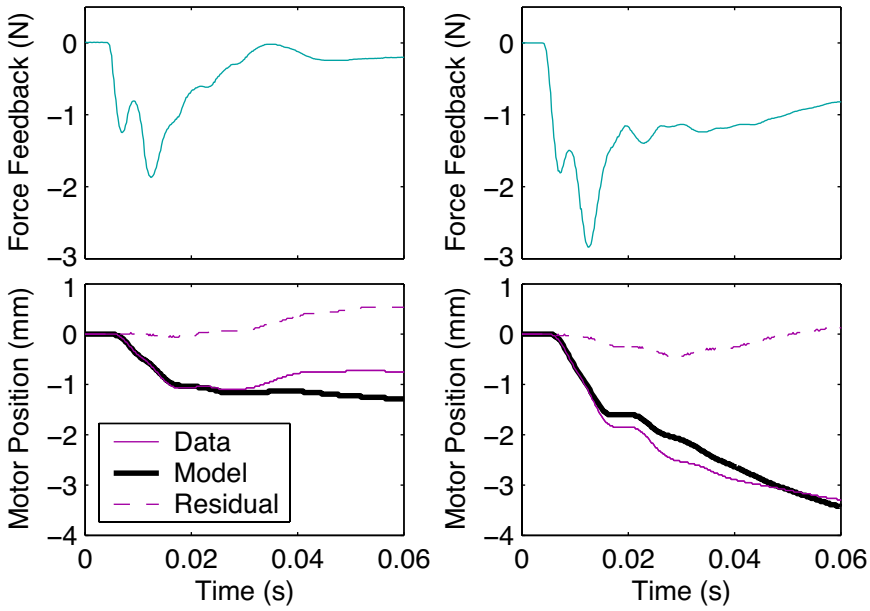

Figure 3.25: Two sample time responses for the real system and the model.

$$m_m\ddot{x}_m = b_c(\dot{x}_d - \dot{x}_m) + k_c(x_d - x_m) + C_m + F_f \tag{3.32}$$

$$m_d\ddot{x}_d = b_c(\dot{x}_m - \dot{x}_d) + k_c(x_m - x_d) - b_d\dot{x}_d - F_{\kappa_l} + C_d \tag{3.33}$$

$$m_h\ddot{x}_h = F_{\kappa_l} - b_u\dot{x}_h - k_u x_h \tag{3.34}$$

$$C_m = \begin{cases} -c_m & \text{if } \dot{x}_m > 0 \\ 0 & \text{if } \dot{x}_m = 0 \\ +c_m & \text{if } \dot{x}_m < 0 \end{cases} \qquad C_d = \begin{cases} -c_d & \text{if } \dot{x}_d > 0 \\ 0 & \text{if } \dot{x}_d = 0 \\ +c_d & \text{if } \dot{x}_d < 0 \end{cases} \tag{3.35}$$

With all of its parameters identified, the dynamic model's ability to predict the master's time response during haptic interactions was tested. Several pre-recorded impact forces were displayed to a user via open-loop output. During these tests, the user held the handle of the haptic interface, moving it forward as though tapping on a remote or virtual environment. The force profile was displayed when the user crossed a position threshold, and the resulting master motion, x_m, was recorded. The model was subjected to the same feedback force, F_f, and its prediction of induced motor motion was recorded.

Two sample test results are shown in Figure 3.25 as experimentally measured, simulated, and residual motor position traces. The entire set of open-loop testing, which was conducted for a variety of transients, users, and incoming velocities, indicates that the full model adeptly predicts system behavior, especially for the first fifty milliseconds after impact before other effects such as the user's hand motion and reflexes alter the response. The full, physically based model can be used to estimate the behavior of the device during operation, and it can also be used to predict the effect of changing various mechanical elements in the system.

3.6 Summary

This chapter explored the internal dynamics of impedance-type haptic interfaces. The combination of computer, device, and user forms a complex electro-mechanical-biomechanical system that exhibits strongly frequency-dependent behavior. As such, simple second-order models fail to capture important high-frequency effects, and more sophisticated modeling methods are needed. The two-step process presented in this chapter enables accurate broad-spectrum characterization of a variety of systems. The multi-element transmission

from motor to stylus makes it well-suited to comprehensive evaluation; this non-parametric frequency-domain technique treats the system as a black box and seeks to identify input-output relationships directly.

If comprehensive evaluation reveals significant nonlinearities or time varying behaviors, successive isolation can be employed to obtain an accurate, high-order, nonlinear model of these dynamics. Successive isolation of system components enables independent identification of Coulomb and viscous friction, inertia, nonlinear and/or hysteretic stiffness, and user impedance. Each of its steps may employ frequency- and/or time-domain testing; using both in combination validates findings along the way and helps elucidate the many targeted parameters.

The models that result from both characterization methods do a good job at capturing the detailed high-frequency behavior of a haptic interface with a single degree of freedom. These techniques can be extended to more complex systems by looking at each joint independently, though the extent to which joint dynamics are coupled will need to be explored. The models obtained in this work can be used to understand the dynamics of individual axes of haptic interfaces and to develop ways to improve the control algorithms used in teleoperation and virtual environments. The relationships from current command to handle acceleration and current command to motor position are particularly interesting, as they can be used to improve the feel of hard contact, as will be presented in the following two chapters.

Chapter 4

High-Frequency Acceleration Matching

An ideal haptic interface would make remote and virtual interactions feel like direct manipulation, portraying the dynamics of the environment with vivid, realistic haptic stimuli. As discussed in Section 1.5, the primary shortcoming of present interfaces is the absence of high-frequency haptic feedback, which is a direct limitation of closed-loop position-based control. Sections 2.1.4 and 2.2.3 described methods that have been developed to add high-frequency vibrations to remote and virtual interactions respectively, but none of these existing approaches has been widely adopted, perhaps because of their complexity and dependence on hand tuning.

In real manipulation, vibrations above about 30 hertz are caused primarily by hard contact between a hand-held tool and a stiff environment. These cues provide rich information about material and surface properties and also help humans identify transitions between task phases, such as the moment when the screwdriver has seated correctly in the head of the screw. Haptic interfaces would feel significantly more realistic if imbued with these important, information-laden, high-frequency vibrations, preferably achieved with standard hardware and simple, extensible algorithms that do not require hand tuning.

Faced with this challenge, I invented the technique of high-frequency acceleration matching, which I occasionally abbreviate as "HFAM." Fundamental to this approach is the belief that hard contact in remote and virtual haptic interactions will feel most realistic when it creates fingertip accelerations that are identical to those produced during real interactions.

Inspired by the prior efforts in vibration feedback that are explored in Section 4.1, I determined that the transmission dynamics of the haptic interface are the most important factor in accurately generating high-frequency fingertip accelerations. These dynamics depend on the behavior of the haptic interface's current amplifier, its mechanical elements, and the user's hand, and they can be thoroughly characterized using the methods of Chapter 3.

The characterized relationship from current command to handle acceleration can generally be inverted, as explained in Section 4.2. This model inversion enables the controller to compute the motor current command that is necessary to create a specified acceleration signal at the handle of the interface. Section 4.3 details the implementation of high-frequency acceleration matching for teleoperation, and Section 4.4 specifies its application to virtual environments. The user study described in Section 4.5 demonstrates the significant realism improvements enabled by adding HFAM to a hard virtual surface, and Section 4.6 summarizes the primary themes of this research.

4.1 Prior Work in High-Frequency Haptic Feedback

As overviewed in Chapter 2, most teleoperative and virtual environment systems rely entirely on closed-loop position feedback to connect the user to the remote or virtual environment, thereby restricting the haptic feedback to low-frequency signals. One notable exception to this trend is an extended body of work on high-frequency haptic feedback that was started by Robert D. Howe at Harvard University and Mark Cutkosky at Stanford University. Through various collaborations and publications, this thread of research has contributed two main paradigms to the literature, one for teleoperation and the other for virtual environments.

4.1.1 Combined Force and Vibrotactile Feedback

In 1995 Kontarinis and Howe demonstrated that operators of a high-fidelity two-fingered master-slave system [62] could perform a variety of tasks more efficiently and more effectively when vibrotactile feedback was added to a standard position-force controller [76]. [1] As illustrated in Figure 2.3, their remote robot's fingertips included accelerometers embedded in a soft rubber skin; detected accelerations were amplified and rendered via a high-frequency

[1] Although its title mentions virtual environments, the 1995 Kontarinis and Howe article develops a vibrotactile feedback methodology for teleoperation only. With some modifications, the described methods could also be used in virtual environments.

vibration display on each finger of the haptic interface. These actuators were comprised of small loudspeakers mounted "upside-down" so that the large base mass was free to move, thereby imposing inertial reaction forces on the user's finger. The authors delineated and tested three types of tasks: those for which detecting high-frequency vibrations is the primary goal, those in which vibrations can improve performance by signaling transitions in task state, and those for which vibrations increase the realism of the interface but do not augment operator performance.

As anticipated, subjects were able to identify a damaged bearing and puncture a thin membrane more accurately when vibrations were added to the system's low-frequency (less than 30 hertz) force feedback. Although combined force and vibrotactile feedback did not enable them to achieve faster completion times for a peg-in-hole task, users commented that the vibrations improved the "feel" of the interface. Considering the significant performance improvements their methods achieved, and the updated vibration actuator this group later developed [30], it is somewhat surprising that other researchers have not added similar vibrotactile feedback to their teleoperation systems, at least to my knowledge. Although the phenomenon is complex, I hypothesize that there are four main reasons that this paradigm has not yet been widely adopted.

1. *Secondary Actuation:* The first factor that may have contributed to the lack of adoption of this technique is that it requires a secondary actuator. Input and output channels are often limited on haptic interface systems, and researchers usually use existing hardware rather than building their own. Augmenting a haptic interface with vibration actuators that are appropriately mounted and tuned requires time, money, and expertise, any one of which may be in short supply for most projects.

2. *Accelerometer Availability:* Adding vibration feedback to a teleoperation system requires a small, robust accelerometer on the end effector of the slave. Kontarinis and Howe used a relatively expensive piezoelectric accelerometer [76], the cost of which may have discouraged adoption. Other sensor types were available at the time [63], but they were not widespread. The first surface micro-machined accelerometer was introduced by Analog Devices, Inc. in 1991, and it was not available in volume production until 1993 [155], approximately the time at which this research was being conducted. The recent development of a wide variety of inexpensive MEMS-based accelerometers has since removed this barrier.

3. *Underlying Controller:* This haptic feedback technique was demonstrated in combination with a high-fidelity position-force controller that had an innate force-feedback bandwidth of 80 hertz. The authors low-pass filtered this feedback to deteriorate its benefits, but the remaining 30 hertz bandwidth is higher than that typically attained with the very common position-position controller. The paper consequently shows only small user performance improvements with the addition of vibrotactile feedback; these improvements would be more significant, applicable, and compelling if the technique were instead showcased alongside position-position control. A similar high-frequency vibration feedback channel was employed on an underwater manipulation system with no other feedback in [30], but it did not include tests of human subject performance with and without the vibrations.

4. *Difficulty of Conveying "Feel":* The final possible limitation to the spread of this technique is one endemic to haptic feedback research: readers of a paper cannot feel the difference between various approaches, and so they may find it difficult to grasp the relative importance of performance increases and user comments. Showing videos or system demonstrations at conferences may help spread successful techniques, but dissemination of such results is often slow.

The promising work performed by Kontarinis and Howe in 1995 is a primary motivator for this research, which seeks overcome the above list of limitations. One element of this previous work that I have also sought to explore further is the fidelity with which fingertip accelerations are rendered. The authors performed an acceleration throughput test by moving the remote robot's fingertip with a vibration test shaker and measuring the accelerations produced at the haptic interface's handle. They found that "the measured response amplitude was flat to within 7 dB across the frequency range of interest." No attempt was made to compensate for these variations, which amount to a magnitude ratio (maximum / minimum) of 2.24, nor to quantify variations in phase at different frequencies. Furthermore, it is not clear from the text whether a user was holding on to the master during these tests; if one was not, the results are less relevant to the performance achievable by the device during use. I hypothesize that even more striking performance benefits could be achieved by more carefully controlling the amplitude and phase of the rendered accelerations at the user's fingertips.

4.1.2 Reality-Based Vibration Transients

The second branch of prior work upon which this research is founded also grew from Howe and Cutkosky's investigations, led by Allison M. Okamura in conjunction with several other colleagues. This group worked to add high-frequency vibrations to hard virtual surfaces through the use of reality-based transients. In 1995 Wellman and Howe first showed that the force response of a stylus tapped on a rigid surface can be approximated as an exponentially decaying sinusoid with a frequency that depends on material composition [152]. Initial human-subject testing demonstrated that the addition of such transients via a secondary voice-coil actuator improved the perception of virtual stiffness during tapping over position feedback alone.

The vibration transient approach was then extended to standard haptic devices by Okamura, Dennerlein, and Howe in 1998, as they superimposed position-based feedback with short, contact-triggered transients via a haptic interface's existing DC motors [114]. The accelerations caused by tapping on a variety of materials were recorded, parametrized, and collected into a library. An arbitrary scale factor between recorded acceleration and transient amplitude was determined by the experimenter and used in human subject tests conducted with an Impulse Engine joystick. This purely algorithmic method was shown to improve the perceived stiffness of virtual objects and to enable more accurate differentiation between various virtual materials.

This project was refined and expanded to three-dimensional exploration in 2000 by Okamura, Hage, Cutkosky, and Dennerlein [113, 115]. The accelerations recorded from contact with many different materials were determined to be too high in frequency for the selected haptic interface, an Immersion 3GM device. The authors concluded that it was thus not viable to exactly replicate the recorded accelerations, and so they scaled them linearly down in frequency to the renderable range of their device. Although they recognized that one could compensate for the system's high-frequency roll-off in software, the researchers choose not to do so for unknown reasons that were most likely based on hardware limitations. Instead of employing dynamic compensation, the authors employed extensive psychophysical testing to determine which vibration parameters felt most realistic for each target material. During this tuning phase, subjects used the haptic interface to tap back and forth between the real surface and the virtual rendering to try to match their respective haptic impressions as closely as possible. A follow-on performance test found that subjects could adeptly distinguish between virtual objects with different parameters.

Though some thought they felt "active" or "not exactly realistic," most users preferred the surfaces with vibrations to those without [113].

What can be learned from this second body of research is that high-frequency transients added at contact can significantly increase the realism of hard virtual objects. In light of this finding, it is surprising that other researchers and companies have not added such feedback to their interfaces, a phenomenon that I attribute primarily to the difficulty of conveying "feel" through conference and journal publications. Unanswered questions on this topic include how to create high-frequency handle accelerations that are beyond the natural bandwidth of the interface, and how to select transient parameters without extensive psychophysical tuning; these questions have strongly motivated and driven my own investigations on this subject.

4.2 Compensating for Haptic Interface Dynamics

Other researchers have developed some methods for displaying high-frequency acceleration transients during remote and virtual interactions, but none has been widely adopted, and open questions about optimal implementation techniques still remain. In response, I have developed the approach of high-frequency acceleration matching to continue this line of research and advance haptic display toward the ever-present goal of portraying realistic hard contact with impedance-type devices.

4.2.1 Primary Goals of HFAM

Building on prior work in this field, I have defined the following five main objectives for the approach of high-frequency acceleration matching:

- During teleoperation, match the user's high-frequency fingertip accelerations to those experienced by the end effector of the slave robot without intermediate scaling or other filtering. This design goal allows the user to feel exactly what the remote robot is experiencing, creating the impression that the master handle is rigidly attached to the slave tip and therefore linking the user more closely with the environment and the task at hand.

- For virtual environments, match the user's high-frequency fingertip accelerations to

those that would be experienced during the real interaction being simulated; for example, record contact accelerations from a wooden stylus tapping on the hood of a real car in order to build a haptic simulation of this interaction. Aligning the acceleration cues will make the two environments feel exceptionally similar.

- Use the haptic interface's existing DC motors to create the desired high-frequency vibrations, rather than adding a secondary actuator to the handle. This choice facilitates implementation on most existing hardware, avoids adding extra mass to the device, and can potentially achieve identical results, assuming the device has a relatively good connection from motor to handle and sufficient actuator power to render the target interactions.

- Compensate for the motor-current-to-handle-acceleration transmission dynamics of the user-device system through careful characterization and model inversion. These analyses permit accurate generation of a wide range of acceleration transients. Once the dynamics of a certain device are identified, its inverted model can be used to make contact with almost any remote or virtual hard surface feel real. Although inverting these internal dynamics cannot enable a system to create high-frequency accelerations that are fundamentally beyond its capabilities, perhaps due to backlash or excessively low bandwidth, this software-based method can significantly extend the high-frequency range of most existing hardware.

- Evaluate the realism of various rendering algorithms by having users perform direct, blind comparisons between virtual and real objects. If a haptic interface is to be used to defuse bombs or train heart surgeons, it must be able to render interactions with a fidelity that is currently unattainable; I believe that we will reach this goal only by challenging ourselves to fool human users, which can be viewed as a Turing test for haptic interaction. [2]

[2] In 1950 Alan M. Turing wrote a seminal paper on artificial intelligence (AI) in which he examined whether a computer could fool a human into believing that it was another human through written interaction alone [148]. The test that I propose is similar in that it seeks to establish whether a haptic interface could fool a human into believing it was a real environment through physical interaction alone. I hope that the haptic Turing test may motivate haptics research in the same way that the original challenge spurred the progress of AI, though this new challenge may prove to be even more difficult due to the remarkable human sensitivity to haptic stimuli.

4.2.2 Inverting a Haptic Interface's Dynamic Model

The foremost engineering challenges for high-frequency acceleration matching are to obtain a good model of the interface's internal dynamics and to devise a reliable method for inverting this model. The first of these two challenges can be addressed with the modeling methods of Chapter 3, and this section presents an approach for inverting such models.

This task is most easily achievable when the user-device system behaves in a linear, time-invariant manner. In this case comprehensive evaluation will yield an LTI model of the relationship between commanded motor current and measured handle acceleration while the device is being held by a user. The system's handle acceleration transfer function is defined by

$$H_{ha}(s) = \frac{A_h(s)}{I_m(s)}, \tag{4.1}$$

with $A_h(s)$ being the Laplace transform of the handle acceleration and $I_m(s)$ being the Laplace transform of the motor current command; the transfer function for the corresponding model is denoted $\hat{H}_{ha}(s)$. A short constant time delay can be characterized via comprehensive evaluation and discarded during this inversion process, as human reaction time to tactile stimuli is about 140 milliseconds [92], and delays less than this duration are not noticeable to the operator [111]. As presented in Section 3.4.6, a dependence on grip force can often be modeled by creating separate LTI models for different grip force levels, including the zero grip force case of the device's natural dynamics, and interpolating appropriately.

If the system does not behave in a linear, time-invariant manner, successive isolation can be employed to isolate the nonlinear and/or time-varying elements. This analysis may uncover amplifier cross-over distortion, a soft hysteretic compliance, backlash in a mechanical joint, or large Coulomb friction. When possible, these components should be fixed or replaced to enable the use of a linear model. If a strongly nonlinear model is necessary, it can be linearized around a certain operating point, or nonlinear model inversion techniques can be utilized. The remainder of this section will deal only with the case in which a linear model has been obtained.

To facilitate real-time inversion, the LTI model should have a relative degree of zero, i.e. it should have the same number of poles as zeros. Physical systems almost always display low-pass behavior, so extra zeros will need to be added just above the frequency range of interest to achieve this effect, adjusting the rest of the model to match the data accordingly. For example, the Phantom model pictured in Figure 4.1 naturally has a relative

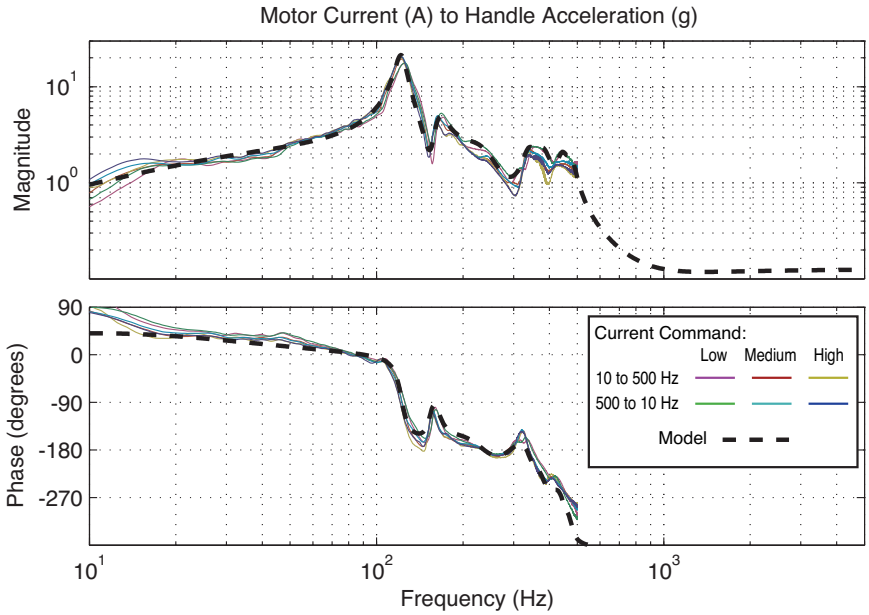

Figure 4.1: Phantom handle acceleration ETFEs and model, $\hat{H}_{ha}(s)$.

degree of two, so a pair of zeros at 800 hertz with a damping ratio of 0.5 were added, for a total of 15 poles and 15 zeros. In the general case, the additional zeros should be placed above the studied frequency range but should be significantly slower than the sampling frequency to ensure adequate discretization. The locations and damping ratios of the other high-frequency poles and zeros should be adjusted along with the placement of the extra zeros to maintain agreement between experiment and simulation in the frequency domain of interest. Having a relative degree of zero gives the forward model a finite gain at high frequency, as seen in Figure 4.1.

Another way to view these added zeros is that they become a low-pass filter (high-frequency poles) when the model is inverted. Selecting their placement in tandem with the model, as suggested here, ensures that this filter does not change the treatment of signals near the identified upper frequency limit of the system model. Selecting the filter's characteristics in light of its net effect on the model facilitates better model inversion than performing these two steps separately.

Once it has an equivalent number of poles and zeros, the linear model for these system dynamics is easily inverted:

$$\hat{H}_{ha}(s)^{-1} = \frac{1}{\hat{H}_{ha}(s)} \tag{4.2}$$

The poles of the model become the zeros of the inverse model, and vice versa. With this inverse model in hand, one can estimate the motor current necessary to create a desired handle acceleration as

$$\hat{I}_m(s) = \frac{1}{\hat{H}_{ha}(s)} A_{h,des}(s). \tag{4.3}$$

Because the modified forward model has a finite gain at high frequency, so does the inverse model; the forward transfer function's added high-frequency zeros simplify the inversion process and implicitly act as low-pass filters in the inverse to limit the amplification of high-frequency noise in the acceleration signals being matched.

To further protect against noise amplification, the desired handle acceleration signal ($a_{h,des}$) should be smooth, whether it stems from a real-time slave-tip reading or pre-recorded hand-held testing. Care should be taken to minimize electrical noise on the accelerometer input lines by using a clean power source and shielded cables. After low-pass filtering to prevent aliasing, the electrical signals should be buffered in analog circuitry before the analog-to-digital conversion, and unused analog inputs should be tied to a stable voltage.

If these measures do not result in an acceleration signal that is sufficiently clean, the signal should be smoothed to remove content above the characterized frequency range of the model before processing with its inverse. The technique of windowing works well for this purpose, but it also incurs a short, constant time delay equal to half of the length of the window. Window durations that are on the order of a few milliseconds can smooth a ten kilohertz signal with a bandwidth in the hundreds of hertz; the high-frequency feedback channel is thus slightly delayed from the closed-loop position loop. This delay does not generally pose a problem because it is small and it naturally combines with the master's output delay upon display.

4.3 Implementing HFAM for Teleoperation

Seeking to provide the user with an authentic feel for the objects contacted by the slave, this section develops a new paradigm for telerobotic feedback that accurately transmits the fine vibratory details of contact. The approach of high-frequency acceleration matching augments a standard bilateral position controller with a feedback channel from the slave tip to the master handle, as illustrated in Figure 4.2. The high-frequency end-effector accelerations that stem from contact with hard or textured objects are measured in real

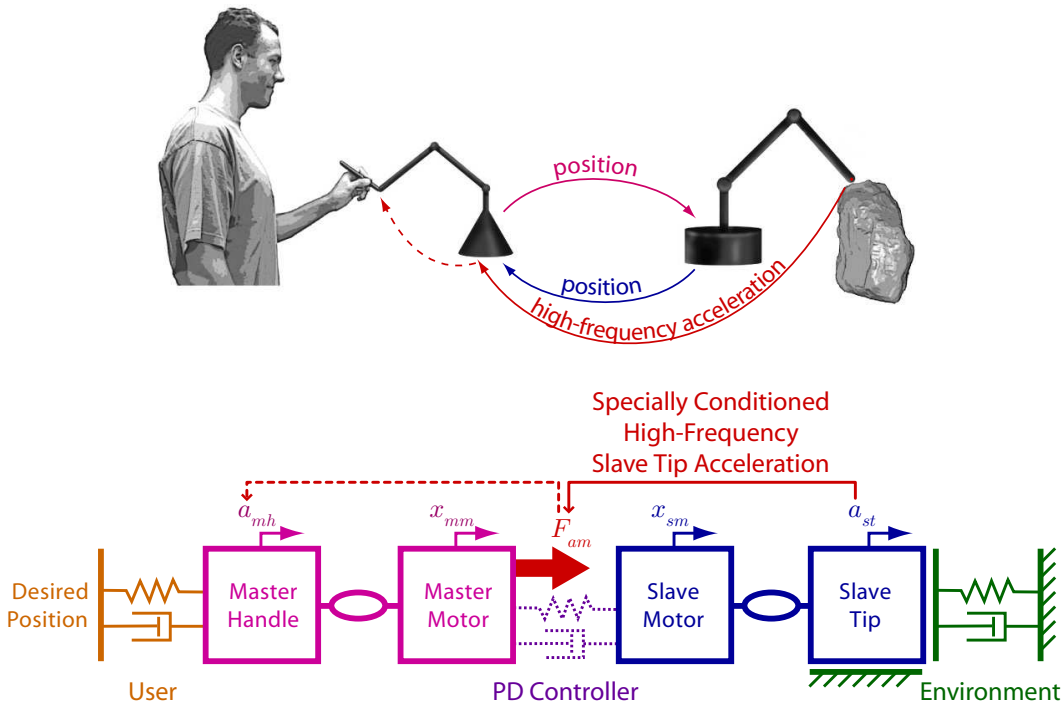

Figure 4.2: Illustration and single-axis lumped-parameter model of a telerobotic system under position-position control augmented by high-frequency acceleration matching.

time and recreated at the user's fingertips. While such an effect can be achieved by adding a secondary actuator to the handle of the master mechanism [30, 76], the device's main motors serve this purpose beautifully if the dynamic connection from motor to handle is taken into account, as first presented in [82]. This section explores the technical details of the HFAM approach and demonstrates its ability to render hard and rough surfaces on a standard master-slave system.

4.3.1 Algorithm Definition

High-frequency acceleration matching provides a new method for connecting the user to the remote environment during teleoperation. The HFAM-augmented controller combines a low-frequency power band with a high-frequency information band, divided at approximately 20 hertz. This approach is inspired by the human's asymmetric sensation and actuation bandwidths and aligns well with Daniel and McAree's frequency separation of human manipulation [28]. As illustrated in Figure 4.2, the method of acceleration matching

augments position-position control with a secondary feedback channel based on the slave tip acceleration, creating a hybrid controller that is better suited to human sensing capabilities.

Although PD control alone cannot convey high-frequency accelerations to the user's fingertips, it adeptly handles quasi-static feedback. Typical hand motions such as tapping and stroking are slower than 10 hertz and are communicated between the sites via bilateral PD control using the motor position signals x_{mm} and x_{sm}. The high-frequency dynamic response of the environment, which generally contains frequencies from several hundred hertz to over one kilohertz, is measured via the acceleration of the slave tip, a_{st}. When humans interact with objects through a mechanical tool, accelerations such as these travel along the structure of the tool to strongly stimulate the Pacinian corpuscles in the human fingertips [12, 154], providing rich information about the object's geometry, texture, and material composition. HFAM for teleoperation applies an additional force, F_{am}, at the master motor to attempt to match the high-frequency acceleration of the master's handle, $\tilde{a}_{mh}$, to that of the slave tip, such that

$$\tilde{a}_{mh}(t) = \tilde{a}_{st}(t), \tag{4.4}$$

where the tilde signifies a high-pass-filtered signal.

Creating specific master handle accelerations by requesting motor current requires knowledge of the dynamics of the user's hand and the master itself, as discussed in Section 4.2. With a well-identified linear model of the relationship between commanded motor current and handle acceleration, $\hat{H}_{ha}(s)$, the teleoperation controller can perform high-frequency acceleration matching. The necessary signal processing architecture is shown in Figure 4.3 along with sample signals.

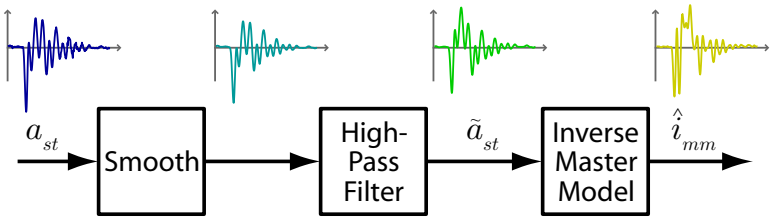

Figure 4.3: Real-time signal processing of slave acceleration.

Smoothing

As the inverse master model typically has large gain at high frequency, the slave acceleration signal, a_{st}, is first smoothed to rid it of any very high-frequency electrical noise or other signals beyond the characterized range of the haptic interface. Smoothing must be performed in real time using an appropriately sized window; as long as the short time delay incurred by this smoothing is on the order of a few milliseconds, it will not affect the user's perception of the interaction, though it may affect the system's stability. The experimental system on which this approach has been developed uses a fifteen-point modified Bartlett-Hanning window, which adds a fixed delay of seven time steps but preserves the important shape of high-frequency transients.

High-Pass Filtering

As shown in Figure 4.3, the smoothed acceleration signal is next high-pass filtered to prevent overlap with the low-frequency power band. The PD controller is responsible for transmitting movement below the tracking bandwidth of the slave's sub-system, and the acceleration-matching channel should not interfere. Furthermore, the high-frequency feedback channel should not attempt to recreate free-space accelerations that stem from user movement, so this high-pass filter should always be set above ten hertz. The experimental system on which these techniques were developed uses a second-order linear filter with a bandwidth of 22 hertz, well matched to the second-order low-pass behavior of the slave's PD controller and significantly above the range of human intention. The discrete-time equivalent of this filter is applied in real-time to the system's smooth slave acceleration signal, producing the high-frequency version, $\tilde{a}_{st}$, that needs to be replicated on the master.

Model Inversion

To determine the necessary master motor current request, $\hat{i}_{mm}$, the smoothed, high-pass-filtered version of the slave tip acceleration is then applied to the inverse of the identified master model, via

$$\hat{I}_{mm}(s) = \frac{1}{\hat{H}_{ha}(s)} \tilde{A}_{st}(s). \tag{4.5}$$

In order to perform this transformation in real time, the inverse model must be discretized at the system's servo rate, using a Tustin approximation to preserve stability. The frequency response of the discretized model should be checked against the continuous-time model; if

the two differ significantly, the sampling rate should be increased. The resulting model is then included in the real-time controller, using double-precision floating point calculations to avoid numerical instability. With such a process in place, a telerobotic controller can determine the current command necessary to match the master handle's high-frequency accelerations to those of the slave tip.

4.3.2 Experimental Results

This section investigates the effect of matching high-frequency accelerations in real time, showing results from tapping on a hard object and stroking a rough texture. The master-slave system on which this research was conducted consists of a pair of early Phantom robots by SensAble Technologies. As pictured in Figure 4.4, each device has three degrees of freedom, incorporating high-fidelity Maxon motors, smooth cable drives, and relatively stiff linkage elements. The orientation of the distal link of the master robot was reversed to point upward and was rigidly extended with a pen-based stylus. Each joint was connected to the corresponding joint on the other device via PD control.

This experimental work focused on adding high-frequency acceleration matching to the vertical axis of the haptic interface, which runs along the length of the handle and bears primary responsibility for tapping feedback. Acceleration matching could be added to the other joints through replication of the current work. The shoulder and elbow joints of the master mechanism were driven with high-bandwidth linear amplifiers taken from an Immersion Impulse Engine 2000, and the system's other four axes were driven by the Phantom's standard pulse-width modulation amplifiers. The linear amplifiers provide excellent high-frequency response, producing full-scale sinusoidal current at up to one kilohertz with no attenuation or phase lag. One drawback of these amplifiers, however, is their 1.4 amp maximum current; although the Phantom's motors can sustain only 1.26 amps continuously, they can tolerate much higher current for short durations, which would allow for even stronger high-frequency haptic cues.

Master and slave accelerations were measured by rigidly attaching a MEMS accelerometer to the endpoint of each device. As shown mounted on a custom printed circuit board in the inset of Figure 4.4, the ADXL321 chip from Analog Devices provides a range of ± 18 g, an adjustable bandwidth that was set at one kilohertz, and a footprint of just 16 square millimeters. The voltage outputs of the two accelerometers were measured using a National Instruments PCI-1200 card, and the system was controlled at a five kilohertz servo

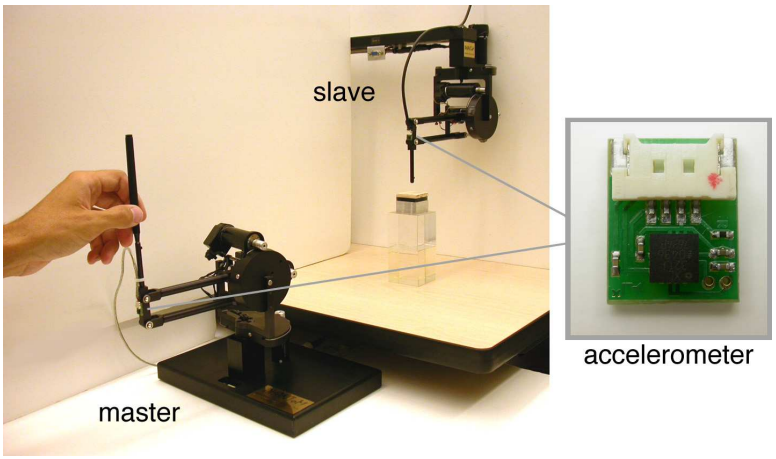

Figure 4.4: Telerobotic testbed.

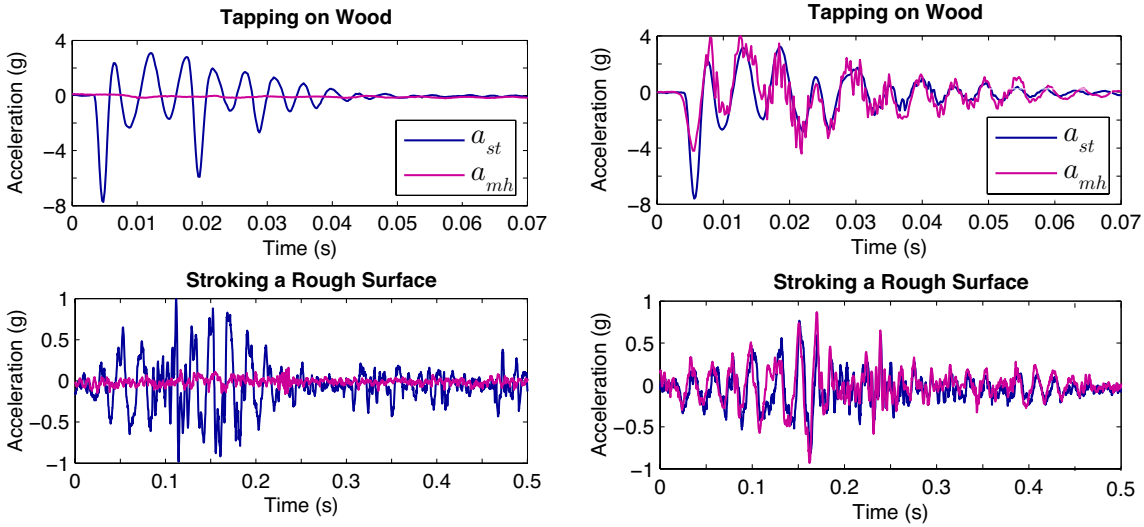

(a) PD Control (b) PD Control with HFAM

Figure 4.5: Slave tip and master handle accelerations for (a) bilateral PD control alone and (b) bilateral PD Control with high-frequency acceleration matching.

rate by a desktop computer running RTAI Linux. This telerobotic system, with its pair of accelerometers, served as an excellent testbed for the development of the high-frequency acceleration-matching approach to haptic feedback in telerobotics.

High-frequency acceleration matching was added to the shoulder axis of the master-slave system using the techniques described in Section 4.3.1. The tip of the slave is an aluminum cylinder, so contact with hard objects generates very high frequency accelerations. Sample master and slave accelerations for tapping on wood and stroking a rough texture under two different control schemes are shown in Figure 4.5: the left plots show PD control alone, and the right plots show PD control augmented with high-frequency acceleration matching. Note that the intervening time delay of 2.2 milliseconds was removed from the master handle signals to enable visual comparison with the slave tip. In sharp contrast to the poor performance of PD control, the HFAM-enabled system successfully portrays the slave tip's high-frequency accelerations to the user, creating a haptic experience that is significantly more rich than that created by PD feedback alone.

In the acceleration-matching approach, differences between slave and master accelerations stem from smoothing, high-pass filtering, and inaccuracies in the master model. The additional high-frequency noise that is visible on the master handle signals under PD control with HFAM stems from an originally unidentified system resonance near 1200 hertz. This artifact does not significantly affect the feel of the interface, as the buzzing is beyond the tactile sensory bandwidth of the user. It could be eliminated by extending the system model to include this resonance so that model inversion would prevent its excitation during playback.

Although the HFAM feedback algorithm for teleoperation has not yet been evaluated in a formal user study, informal testing indicates that acceleration matching vastly improves the user's ability to determine the material or texture of the environment, as accomplished by the similar strategy and secondary actuator of Kontarinis and Howe [76]. Anecdotally, users have enjoyed the improved sense of telerobotic touch that high-frequency acceleration matching provides, commenting that it makes the system feel more like a rigid tool. Adding high-frequency acceleration matching to systems under bilateral PD control allows the user's fingertips to experience the same high-frequency accelerations as the slave's end effector. Rather than feeling like soft, smooth foam, hard remote objects feel crisp, and textured objects feel rough. The user can take advantage of his or her vast experience with everyday manipulations to interpret these signals and guide the interaction accordingly.

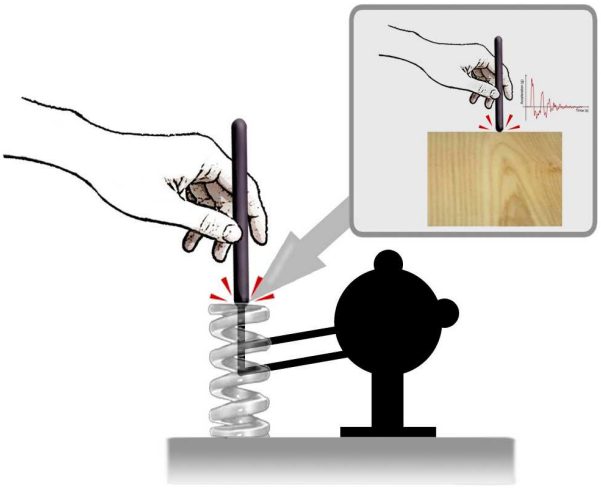

Figure 4.6: High-frequency acceleration matching for virtual environments.

4.4 Implementing HFAM for Virtual Environments

Aiming to make firm virtual surfaces feel indistinguishable from their real counterparts, this section develops a new paradigm for virtual environment feedback that accurately recreates the high-frequency accelerations of hard contact. The approach of high-frequency acceleration matching augments standard position feedback with short surface-specific transients at contact, as illustrated in Figure 4.6. These transients are generated in advance from interactions between a human hand and the object or surface being simulated. The algorithm accounts for the dynamic connection between the haptic interface's motor and handle to accurately recreate the target accelerations at the user's fingertips. This section presents the technical details of this approach and demonstrates its ability to portray hard surfaces on a typical impedance-type device.

4.4.1 Transient Generation

High-frequency acceleration matching provides a new method for generating transients for use in event-based haptics, the virtual environment feedback architecture described in Section 2.2.3. In the past, researchers have used parametric transients such as short-duration pulses and decaying sinusoids, but signal features such as duration, frequency, and nominal magnitude always needed to be hand-tuned for each device and target object, a process that is more art than science. As discussed in Section 4.1.2, extensive psychophysical testing can

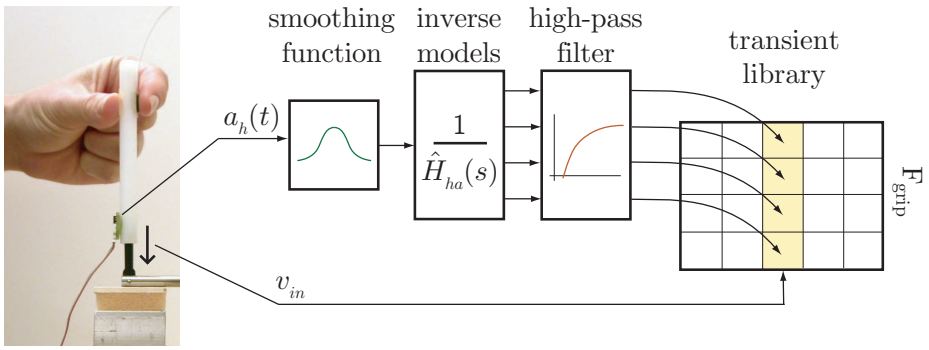

Figure 4.7: HFAM transient library generation.

be used to identify preferred values, but there has not before been a deterministic method for generating transients that feel realistic. Working from a desire to recreate the accelerations caused by real contact, I developed an analytical method for generating a library of current commands that can reproduce pre-recorded handle accelerations. [3]

Recording Accelerations

As diagrammed in Figure 4.7, an instrumented stylus is used to record the contact acceleration transients that result from hand-held contact with the target object. The haptic interface itself can be used for this testing, or another stylus can be used with any other lightweight digitizing robotic arm or even a non-contact position tracking system. Regardless of the position measurement method, a small, high-bandwidth accelerometer should be attached to the stylus. This work has focused on the accelerations that occur along the length of the stylus, as these are most important during perpendicular tapping, but a three-dimensional accelerometer could be used instead.

Impacts are recorded at a wide range of incoming velocities, seeking to adequately cover expected values, as the dynamics of contact depend strongly on this variable; the magnitude of the acceleration transient generally increases with velocity, and its shape often changes as well. Contact accelerations are generally independent of grip force, except at very low levels when the stylus begins slipping out of the human's grasp, so it is not necessary to vary or measure this parameter during the recording phase [38]. For each contact, 100 milliseconds

[3] I have chosen to split the HFAM algorithm into an offline portion, which processes the desired acceleration transients and generates the current waveform library, and a real-time portion, which interpolates the transient from the library based on incoming velocity and possibly grip force. With adequate computational power, the recorded accelerations could be stored instead, and the signal processing could be done in real time.

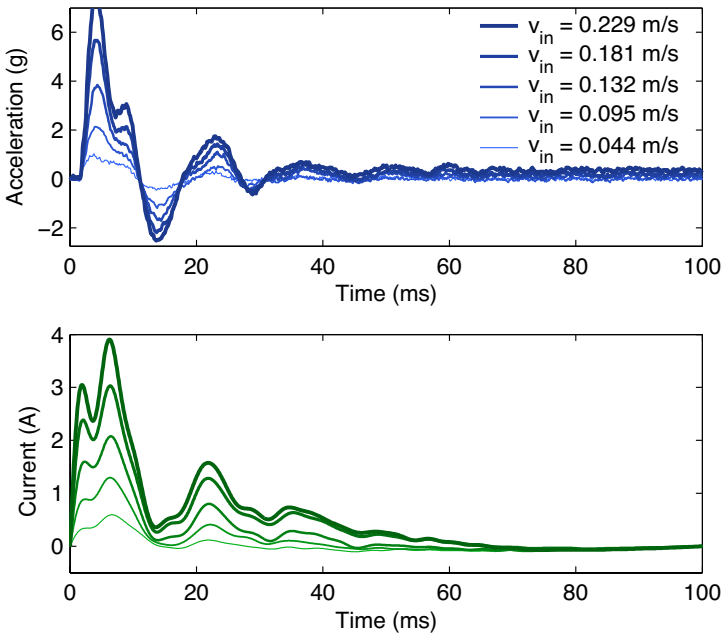

Figure 4.8: Recorded accelerations and current transient library for tapping on a sample of wood on a foam substrate at different velocities using the Phantom.

of acceleration starting just before the impact should be recorded at a high data rate, such as ten kilohertz, and tagged with the incoming velocity. The entire set of recorded acceleration transients should then be plotted together, and a nicely spaced subset should be chosen. Traces that have high incoming acceleration should be discarded, as should ones that appear atypical. A sample set of chosen accelerations is shown in the top half of Figure 4.8; this set can be viewed as the haptic impression of the target interaction.

Smoothing

The selected acceleration signals must be processed and applied to the haptic interface's inverse model to determine the current commands necessary to create them on the haptic interface. Any very high-frequency noise in the acceleration transients will be amplified by the inverse model, since it has finite rather than zero gain at infinite frequency. Each transient is thus pre-padded with zeros and smoothed to remove high-frequency electrical noise without altering phase. As with high-frequency acceleration matching in teleoperation, a variety of window shapes can be used, and the bandwidth of this operation should be set no higher than the one kilohertz bandwidth of human perception.

Model Inversion

The smooth acceleration transients are then applied to the inverse of the system model, producing a raw version of the required current transients. The equation governing this transformation is

$$\hat{I}_{mm}(s) = \frac{1}{\hat{H}_{ha}(s)} A_{h,tr}(s), \tag{4.6}$$

where $\hat{I}_{mm}(s)$ is the Laplace transform of the estimated master motor current command and $A_{h,tr}(s)$ is the Laplace transform of the handle acceleration transient. In certain situations, the identified system model consists of a set of LTI models rather than a single model; Figure 4.7 depicts this occurrence for grip-force compensation under comprehensive evaluation, which was described in Section 3.4.6. In such cases, this transformation should be performed independently for each model on each transient in order to determine the array of current transients necessary to produce that acceleration for all characterized levels of grip force.

High-Pass Filtering and Tapering

The presence of a human user prevents the inverse model for a haptic interface's dynamics from being accurate at low frequency. Humans can actively control the motion of their hands below eight hertz, so the controller should not attempt to generate accelerations in this frequency range. Additionally, the short duration of the transient cannot be used to present such low-frequency feedback, so it should be eliminated from the current transients. A second-order high-pass filter with a cut-off near ten hertz works well for this purpose. The zero padding is subsequently removed and the tails of the signals are tapered to zero to smoothly transfer feedback responsibility from the open-loop transient to the proportional controller. For each system model, the result is a library of current transients like that shown in the bottom plot of Figure 4.8. Note the non-linear variation of both acceleration and current with velocity, which highlights the ability of a model-based approach to produce realistic, non-parametric accelerations at impact.

Verification

When a library is created, the model inversion process should be verified by applying the computed current transient to the forward model and by testing the transients on the

actual hardware. The output from each trial is compared with the desired acceleration transient for a variety of incoming velocities to ensure that the inversion process does not distort the signals significantly. Close correspondence can be achieved with the above methodology, supporting the viability of model inversion for matching virtual feedback to real accelerations. Once this process has been established, transient libraries for a wide variety of materials and objects can be created very quickly. The only hand-controlled step is that of transient selection, which could be automated fairly easily.

In my experience, the above signal processing methods produce transients that feel very realistic to most users when overlaid with a wide range of underlying proportional feedback stiffnesses. Interestingly, though, a small percentage of users will report that the HFAM surface feels slightly active. I hypothesize that these individuals present a somewhat lower hand impedance than other test subjects, so the handle moves more and creates fingertip accelerations that are larger than intended. People are very sensitive to this over-stimulation, as it is not a sensation encountered during interactions with real objects. This impression can be corrected by multiplying the library transients by a gain of 0.8 or 0.9 for these users, rather than portraying them at their full magnitudes. A more thorough approach is to re-characterize the dynamics of the interface for users with low hand impedance and generate a separate library for their use.

4.4.2 Transient Overlay

Once the transient library for a target surface has been created, its high-frequency acceleration-matched transients are added to the standard proportional feedback controller as shown in Figure 4.9. The real-time controller loads the library as a two- or three- dimensional array: the two axes of time (not shown) and incoming velocity will always be present, and a third axis may be used on systems whose dynamics vary strongly with grip force or another parameter. The program also loads the geometry of the simulated environment and performs collision detection between it and the user's proxy. When the user enters a virtual object, a threshold-crossing event occurs, and the controller computes the appropriate open-loop signal from its transient library.

The real-time controller uses measurements of incoming velocity and grip force to construct the appropriate transient from the library. The user's incoming velocity is compared to those of the library transients, and a linear combination of the closest signals is selected. For example, if the incoming velocity was 0.07 m/s, our library's 0.065 m/s and

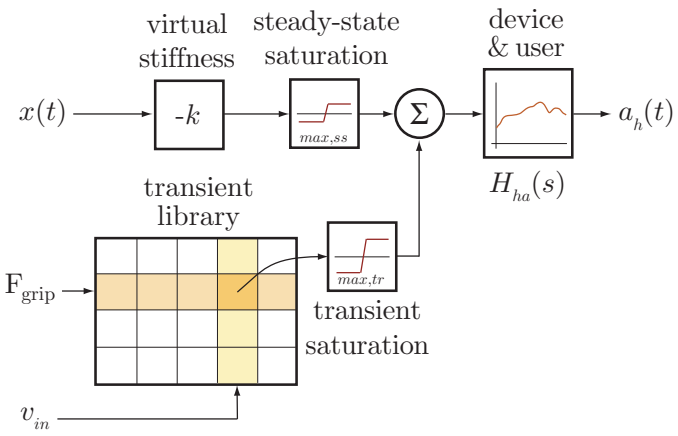

Figure 4.9: Algorithm for combining proportional feedback with grip-modulated high-frequency acceleration-matched transients.

0.081 m/s traces in Figure 4.8 would bracket the current state, and the system would select 0.3125 times the lower transient plus 0.6875 times the upper transient. Linear interpolation is motivated by the commonly observed linear relationship between transient acceleration magnitude and incoming velocity, as demonstrated for example in [113]. If the contact responses for a certain probe/material combination were found to increase nonlinearly with velocity, the interpolation could instead be performed in a nonlinear manner. If grip force or another time-varying parameter is also used as a library index, its present value is measured at the start of contact, and the neighboring levels are again interpolated appropriately. The final open-loop waveform is constructed as a weighted sum of the neighboring transients, and it is added to the system's proportional feedback for the first 100 milliseconds after contact.

As depicted in Figure 4.9, separate current saturation levels can be used for steady-state and transient feedback: DC motors can easily withstand higher current for short durations, as are needed for event-based transients, provided that the temperature of the motor is accurately estimated in real time and used to terminate or attenuate the simulation's feedback if necessary [38]. Exceeding the standard steady-state current limit increases the maximum usable current magnitude and therefore the acceleration transient frequency that can be rendered with a system. These high current levels are available in most haptic systems but are not leveraged by traditional feedback algorithms, adding to the appeal of an event-based approach with HFAM transients. During an interaction, the combined proportional

and acceleration-matched feedback signals are summed and applied to the user together to create the handle acceleration $a_h(t)$, which strongly resembles the target acceleration when the model inversion process is performed well.

4.4.3 Experimental Results

This section shows representative results from the use of high-frequency acceleration-matched transients in simulating contact with hard surfaces. Testing was performed in two phases. The first took a simple approach and did not compensate for changes in user grip force; this apparatus was based on a Phantom device and was used extensively in the human subject study presented in Section 4.5. The second phase was conducted on a custom one-degree-of-freedom haptic interface and included compensation for changes in user grip force. Results from both of these projects are presented below. In both cases, the target interaction was chosen to be tapping on wood, a familiar material whose characteristic transients fall within the current creation capabilities of the utilized interfaces. The higher-frequency signals caused by contact with stiffer substances such as metal require actuation power that is beyond the capabilities of our present amplifiers, though the same methods can be applied to their creation.

Phantom

The first haptic interface on which this work was conducted is the same early Phantom device that was used as the master in the teleoperation system discussed in Section 4.3.2. As pictured in Figure 4.10, this stiff mechanism uses a high-bandwidth linear amplifier to allow transmission of high-frequency signals to the user's hand. During interactions, motion along the axis of the stylus was isolated using a proportional controller to keep the stylus vertical. The motor on the shoulder joint was used to render all interaction forces, and mechanical stops were added to the base joint to keep it centered. Software gravity compensation allowed the stylus to maintain a constant position when the user was not holding it.

Handle accelerations were measured via an Analog Devices ADXL150, which has a bandwidth of one kilohertz and a range of ± 50 g. The small package of this accelerometer was attached to the Phantom's distal link using double-sided tape, and its wires were routed along the arm. The voltage output of the accelerometer was measured using a National Instruments PCI-1200 card. A desktop computer running RTAI Linux sampled

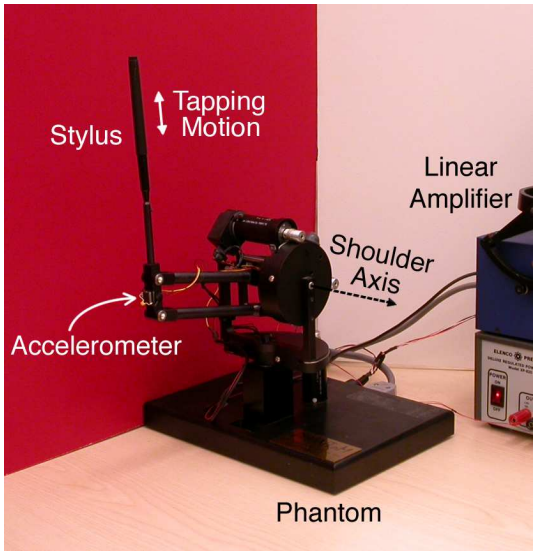

Figure 4.10: Phantom testbed for vertical tapping.

the accelerometer signal, as well as the Phantom's encoders, at ten kilohertz, commanding output from the current amplifier at the same rate. This high servo frequency was chosen to allow the system to measure accelerations and output currents at many hundreds of hertz. Grip force was not sensed, and users were asked to maintain a consistent, moderate hold on the stylus during this informal testing and during the follow-on human-subject study, which is described in Section 4.5.

Sample performance of the acceleration-matching approach to event-based transient generation is shown in Figure 4.11. For the left-most plot, the proportional feedback gain k was set to 680 newtons per meter, the highest value attainable on this system that does

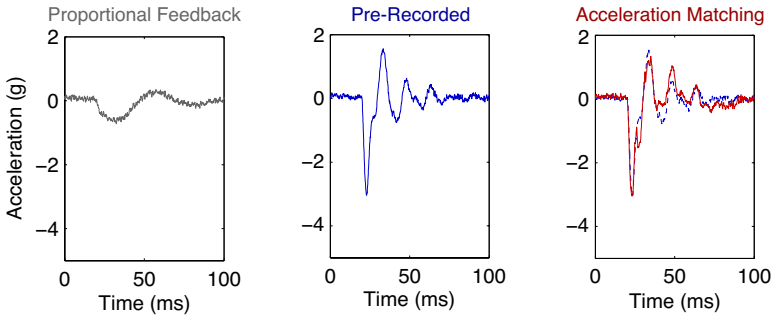

Figure 4.11: Handle acceleration transients at contact under an incoming velocity of 0.11 meters per second.

not create noticeable vibrations at each encoder tick passing. The acceleration transient created when tapping on this surface at an incoming velocity of 0.11 meters per second is seen to be much smaller in magnitude and slower in frequency than the target acceleration. The pre-recorded acceleration transient stemmed from interactions with a piece of wood on a foam substrate, the target interaction for the library shown in Figure 4.8. When the high-frequency acceleration-matched transient is added to the proportional feedback, the user experiences contact accelerations like those shown in the rightmost plot of Figure 4.11: the transient strongly resembles the pre-recorded acceleration, showing large magnitude, high-frequency, and a similar decay rate. Testing like this was conducted on the Phantom for several individuals at a range of incoming velocities with similarly good results, providing initial verification for the performance of high-frequency acceleration matching and indicating its readiness for a formal user study.

Custom One-DOF Interface

A second set of experiments was conducted on a different haptic interface to explore the dynamic effect of user grip force during tapping interactions with HFAM. The chosen system is a custom one-degree-of-freedom (one-dof) testbed, as depicted in Figure 4.12. This mechanism has a single DC motor and permits vertical stylus motion only. This platform incorporates a Maxon RE025 motor with an optical encoder and provides approximately the

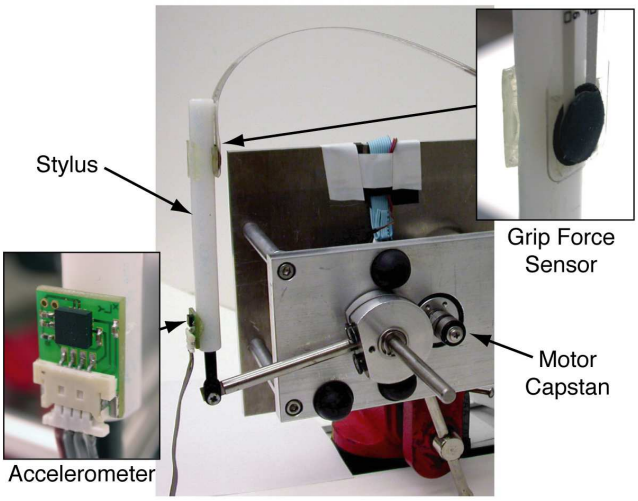

Figure 4.12: Custom one-dof interface for virtual tapping.

same motor-torque-to-stylus-force ratio as a Phantom. The stylus was instrumented with an Analog Devices ADXL321 ±18g accelerometer mounted to a custom circuit board, used for both transient recording and validation. In addition, a FlexiForce A201-1 force-sensitive resistor (FSR) [40] was mounted on the stylus to measure grip force in a two-finger grasp configuration. A National Instruments PCI-1200 card provided analog-to-digital capture resolutions of 0.0175 g and 0.009 newtons for the accelerometer and grip force sensor respectively. The system was controlled by a ten kilohertz servo loop on a desktop computer running RTAI Linux [127].

The low-pass characteristics of this system's PWM current amplifier were characterized and counteracted through the use of a high-pass current-command pre-filter [38]. The remaining dynamics were characterized at discrete grip force levels through comprehensive evaluation, yielding the models shown in Figure 3.12. A transient library was constructed for the wood on foam sample using the methods presented in Section 4.4.1, and it was tested for a range of grip forces and incoming velocities with a single user.

Figure 4.13 shows twelve acceleration transients for virtual taps using grip-modulated acceleration-matched haptic feedback, along with the motor current transients used to create each of them. The black dashed traces in the three left-hand columns are the measured accelerations from real taps that were used to create the transient library. The correspondence between the rendered and desired signals and the uniformity across grip force for each incoming velocity demonstrate that the changing user/stylus dynamics are accurately accounted for by the system. For comparison, if the system were to blindly apply the medium grip current transient regardless of grip force, the acceleration would be too large for softer grips and too weak for firmer grip forces.

Closer examination of Figure 4.13 reveals extraneous high-frequency vibrations for the firm grip force at 30 and 40 cm/s velocity, as well as for the medium grip force at 40 cm/s. At these higher grip forces and higher incoming velocities, the transient current command reached the saturation value of four amps, as can be seen in the right-hand column of the figure. Truncating the transient in this manner distorts the experience of contact and can excite higher-frequency modes in the system in the same way that a step command can. A higher saturation limit could have been used to avoid this effect, though motor thermal limits would have more severely limited the number of strong taps that could be conducted in a short period of time. As with the Phantom interface, the one-dof system achieved good acceleration matching between pre-recorded transients and transients

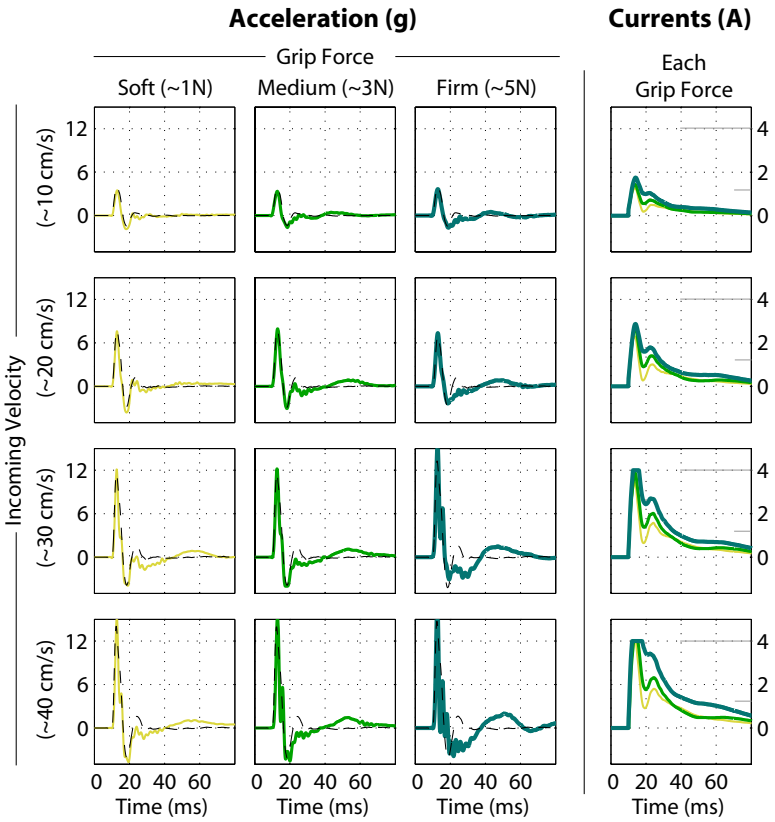

Figure 4.13: Contact accelerations and corresponding current transients for a range of incoming velocities and grip forces.

generated during virtual tapping when saturation was not encountered. The effect these accelerations have on the user's perception of the interaction is explored and discussed in the following section.

4.5 User Evaluation

Realism of virtual contact is inherently difficult to quantify and can be accurately assessed only by conducting perceptual user tests; consequently, the method of high-frequency acceleration matching for virtual environments was evaluated in two carefully controlled user studies. The first investigation performed a preliminary examination of user perception of realism during contact with real and virtual wood, and the results were reported by Kuchenbecker, Fiene, and Niemeyer in 2005 [78]. The experimental methods of this study

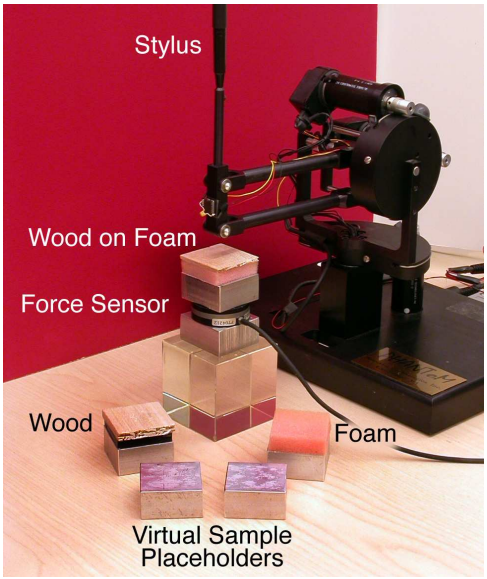

Figure 4.14: Phantom, instrumented test-sample platform, real samples, and virtual sample placeholders.

were refined, the revised experiment was performed with a new set of human subjects, and the results were published by the same set of authors in 2006 [79]. This section presents the methods and findings of the second study, showing that event-based transients, including those designed through high-frequency acceleration matching, can portray hard contact with significantly higher realism than the traditional haptic feedback architecture of proportional control.

4.5.1 Experimental Setup

User testing was performed on the hardware described in Section 4.4.3 with motion of the stylus constrained to one degree of freedom to isolate tapping motions. Tapping on wood was chosen as the target interaction, as most people are familiar with this material from daily interactions. Real samples and virtual sample placeholders were positioned on a rigid stand beneath the center of the stylus, as shown in Figure 4.14. An ATI Mini-40 force sensor allowed the system to measure contact forces at one kilohertz during interactions with real samples, providing a force resolution of about 0.0013 newtons . Interactions with virtual surfaces were portrayed with the Phantom's shoulder motor.

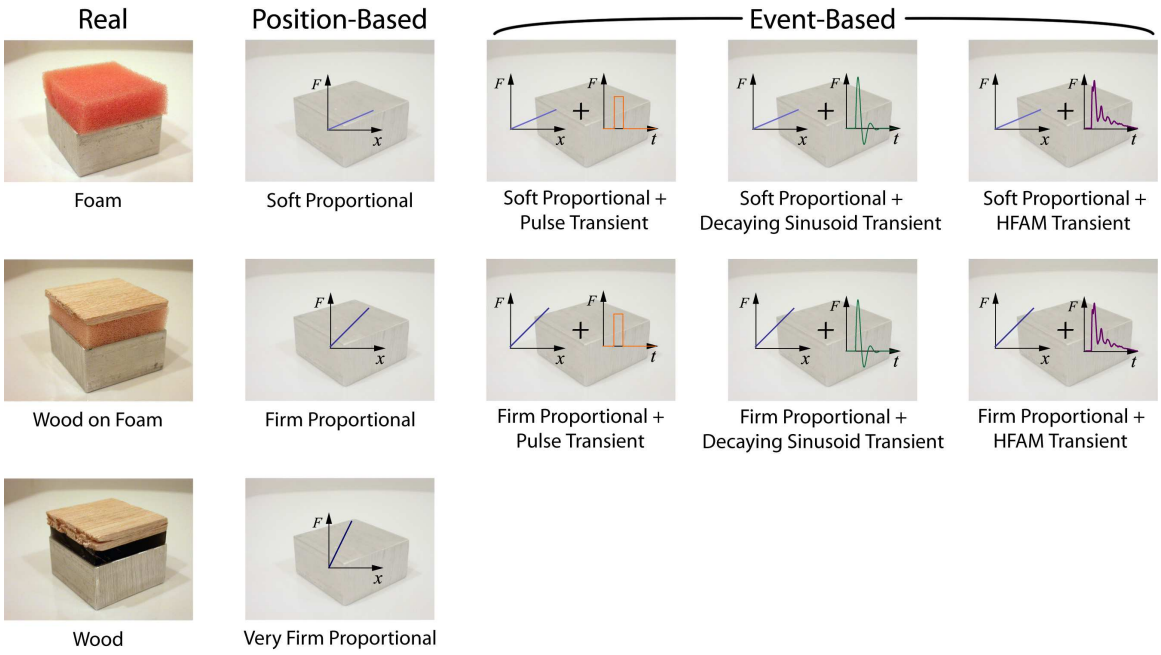

Figure 4.15: Twelve test samples.

Table 4.1: Test sample parameters.

Wood	K $\approx$ 70,000 N/m
Wood on Foam	K $\approx$ 350 N/m
Foam	K $\approx$ 220 N/m
Very Firm Proportional	K = 1020 N/m
Firm Proportional	K = 680 N/m
Soft Proportional	K = 340 N/m
Fixed-Width Pulse	A = 4.55 Ns/m, d = 0.020 s
Decaying Sinusoid	A = 15.9 Ns/m, d = 0.055 s, f = 55 Hz

4.5.2 Test Samples

The feel of the high-frequency acceleration-matched (HFAM) transients developed in Section 4.4 was evaluated in a user study via comparison with three real objects and several other virtual rendering techniques. As depicted in the left-hand column of Figure 4.15, the three real samples were a piece of foam, a piece of wood on top of a foam substrate, and a piece of solid wood. Each material is mounted on an aluminum base to make the top surface a consistent height and to facilitate keyed placement on the stand. The approximate linear stiffness of each of these real samples was measured using the Phantom and the force sensor, as listed in Table 4.1.

The first three virtual samples were designed to represent traditional position-based feedback, as illustrated in the second column of Figure 4.15; the highest proportional gain (very firm) was chosen near the limit of the device's capabilities, at which point encoder discretization causes significant buzzing. The mid-level gain (firm) was tuned to provide strong feedback without encoder buzzing, and the lower gain (soft) was set to half this level. The remaining six virtual samples represent the event-based rendering approach, combining a contact transient with either the firm or soft proportional controller. Overlaying transients on the very firm proportional feedback requires short-duration force output that is beyond the system's capabilities and was thus not included in user testing.

The steady-state stiffness of wood is approximately seventy times greater than the highest proportional gain achievable on the Phantom (see Table 4.1), but only the firm and soft position gains could be used to underlie the event-based feedback. To create contact accelerations that are consistent with this lower underlying stiffness, all of the transients were tuned to duplicate the feel of wood on a foam substrate rather than solid wood. In personal experimentation and informal testing, I have found that this new tuning technique produces virtual feedback that feels qualitatively more real than transients matched to the response of solid materials. Very high-frequency transients feel discordant when overlaid with the relatively soft proportional feedback achievable with today's impedance-type devices, and the lower frequencies generated by hard materials mounted on a soft substrate feel well matched. This phenomenon may explain why Okamura et al. were able to produce realistic-feeling virtual surfaces by scaling down the frequencies of their decaying sinusoid transients, as discussed in Section 4.1.2. Transient generation strategies for velocity-scaled pulses, decaying sinusoids, and acceleration-matched signals are discussed below.

Fixed-Width Pulse

The first investigated transient is the pulse, a simple signal that requires little computation during generation. While some previous work investigated fixed-magnitude, varying-duration pulses [64], this experiment used a pulse of fixed duration d and varying magnitude, scaling the nominal amplitude A by incoming velocity:

$$F_{pulse} = A\,|v_{in}| \quad \text{for} \quad 0 < t \leq d \tag{4.7}$$

This strategy keeps the frequency content of the transient approximately constant over all incoming velocities, as is observed during contact with real objects. The pulse's width and nominal magnitude were tuned by hand while tapping repeatedly on the sample of wood on foam and the virtual surface. The amplitude was further adjusted after some subjects in the preliminary study reported that the pulses felt too active [78]. These and all other chosen transient parameters are given in Table 4.1.

Decaying Sinusoid

The use of decaying sinusoids was also investigated, following the observation that transient accelerations often resemble exponentially decaying sine waves. The nominal magnitude A of this transient is also scaled by incoming velocity:

$$F_{sine} = A\,|v_{in}|\,e^{\ln(0.01)t/d}\sin(2\pi f t) \quad \text{for} \quad 0 < t \leq d \tag{4.8}$$

The frequency f and duration d were selected to match accelerations recorded from tapping on the wood-on-foam sample, picking duration to be an even multiple of the sinusoid's half period for smooth overlay. The nominal magnitude was tuned by hand via comparisons with the specimen of wood on foam.

High-Frequency Acceleration Matching

The final investigated transient type was the high-frequency acceleration-matching approach, which is described in detail in Section 4.4. This method characterizes the complex dynamic relationship between motor current and handle acceleration; the inverse of this model is then used to transform desired acceleration profiles, which may be intricate and difficult to parametrize, into transient current commands. By recording accelerations for contact with the same object under a variety of conditions, a force transient library can be assembled without extensive parameter tuning. Each signal is characteristic of the real situation that produced it, including incoming velocity, hand impedance, and object contact location; for this work, we focused on the dominant effect of incoming velocity, v_{in}, which increases the magnitude and smoothly changes the shape of the transient.

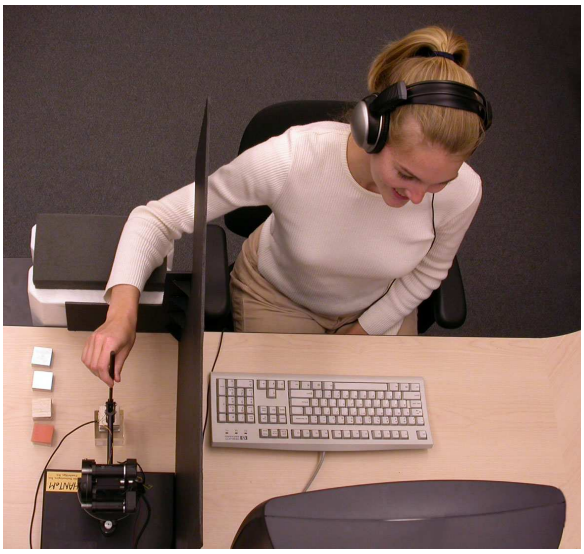

Figure 4.16: Setup for blind tapping on real and virtual wood.

4.5.3 Experimental Procedure

At the beginning of each experiment, the experimenter collected information on the subject's age, gender, handedness, and prior experience with haptic devices. The experimenter then outlined the three phases of the study: familiarization with the wood sample, demonstration of the twelve test samples, and repeated rating of sample realism. Users were told that they would be presented with a number of different renderings of a hard, wooden surface. They were asked to rate, on an scale from one to seven, how well each sample represented the experience of tapping on the real piece of wood. Subjects were asked to repeat the definition of this realism metric before starting the experiment to ensure comprehension.

To isolate the user's sense of touch, extraneous stimuli were removed from the experimental setting, as illustrated in Figure 4.16. Sitting at a computer terminal, the user passed his or her right arm through an opening in a tall barrier that prevented visual observation of the device and samples. The user rested his or her elbow on a padded armrest to counteract muscle fatigue. The user was instructed to hold the stylus with a consistent grasp and to avoid touching the table with his or her left hand to prevent inadvertent transmission of contact vibrations. The user wore headphones playing white noise at a high volume to mask the sounds caused by tapping on the different samples. Simple text commands were

presented on the computer monitor to guide the user through the three phases of the experiment. The operator sat behind the barrier at another computer, monitoring the progress of the experiment and placing samples beneath the stylus.

During the first phase of the experiment, the user was able to tap repeatedly on the real wooden sample to become familiar with its response. The interaction was monitored with the force sensor, and forces known to exceed the amplifier's current limit (had the controller been in use) were detected and indicated to the user by auditory feedback. During virtual tests, the same low tone was provided whenever the commanded force actually exceeded the amplifier's current limit. Throughout the experiment, all trials that evoked this auditory cue were repeated to prevent the system's force limit from biasing the data.

When the user was done interacting with the wooden sample, the system transitioned into the demonstration phase, in which each of the twelve samples, both real and virtual, were presented to the user once in random order. This phase was included to allow the user to learn the experimental procedure, which was replicated in the following testing phase, and to explore the range of samples before beginning to rate their realism. Before each tap, the system would move the stylus to a home position above the sample, giving the operator space to place the next object. Two virtual placeholder blocks were used so that the operator removed and placed an item on the stand every trial, regardless of whether the sample was real or virtual. When the sample was ready, the user was instructed to tap, both by a text command on the monitor and by a recorded voice in the headphones.

The user would then move the stylus down to tap on the surface of the object, which was always at the same height. From the time of first impact, they were given five seconds to tap repeatedly on the surface. After five seconds, the device moved the user's hand back to the home position, and the user was instructed to rate the realism of the sample on a scale from one to seven using the keyboard. The test was repeated if the contact force exceeded the device's capabilities at any time during the tapping, or if the user overpowered the homing force and tapped on the object after the five seconds had expired. Finally, the user could reject a trial by typing an 'x' instead of a digit, indicating that a mistake had occurred such as letting go of the stylus.

Following completion of the demonstration phase, the user proceeded to the testing phase, wherein each sample was presented three times in random order, for a total of 36 trials, plus any repeats for excessive force or rejection. The testing procedure was identical to that of the demonstration phase, and the entire experiment lasted about ten minutes.

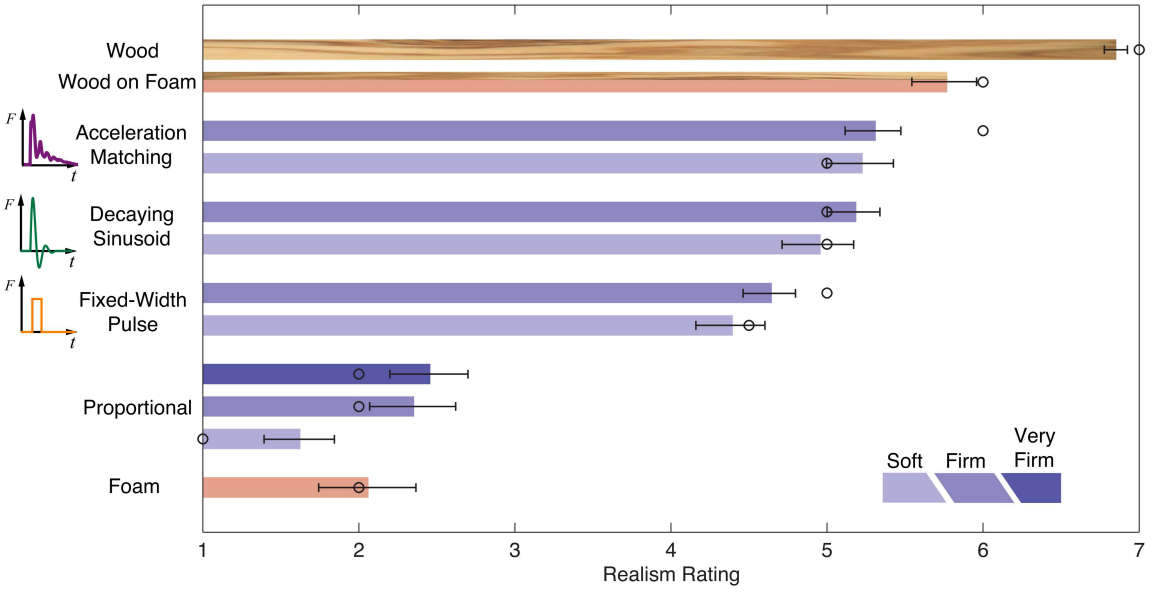

Figure 4.17: Realism ratings of the twelve test samples: bars and circles indicate the mean and median across all tests, and capped lines show the standard error of the sixteen-subject sample.

A short debriefing session followed the completion of the testing, wherein subjects were asked to state the criteria they had used to evaluate sample realism and comment on the experimental procedure.

4.5.4 Results

The user study included sixteen subjects, none of whom had participated in the preliminary version of this study described in [78]. Individuals ranged in age from 19 to 33 and included four females and twelve males. Three of the subjects were left-handed, though all completed the experiment with the right hand. Seven of the subjects had never used a haptic device before. Seven had used such systems a few times, and two reported using haptic interfaces on a regular basis. During the familiarization phase, subjects tapped on the real sample between four and twenty-two times. For each subsequent trial, the system recorded the sample, saturation, rejection, and the set of incoming tap velocities and penetration depths.

The system also stored the user's rating for each trial, indicating the degree to which the user believed the sample felt like real wood. Each of the twelve samples was rated three times by each of the sixteen subjects. Each sample's average realism rating for valid tests,

pooled across users, is shown in Figure 4.17, with higher values indicating higher perceived realism.

Users exceeded the system's force capability an average of six times during the testing phase, with a standard deviation of 8.5. Saturations occurred most frequently for the real wood and the event-based virtual samples. As described above, all such tests were returned to the sample pool to be randomly drawn again and thus are not included in the presented results. The average incoming velocity, pooled across subjects and non-saturated test-phase taps, was 0.10 m/s, with a standard deviation of 0.042 m/s. Within each sample, no significant correlation was found between incoming velocity and realism rating.

4.5.5 Discussion

Both quantitative and qualitative data on the realism of high-frequency acceleration matching was collected during the study. This section analyzes each of these classes of results in turn.

Realism Ratings

Statistically significant differences were found among the ratings given to the twelve samples shown in Figure 4.17, indicating that some samples felt very similar to wood and others were poor imitations. Not surprisingly, the most highly rated sample was real wood, indicating that users maintained a good sense for the response of this common material throughout the testing. Wood was followed interestingly by the wood on foam and the six transient-overlaid virtual surfaces, with acceleration matching rated most highly, though not significantly so. The ratings given to the three proportional controllers were comparable to those assigned to foam, significantly lower than those given to the event-based samples.

To evaluate the significance of these ratings, paired t-tests were conducted on average user ratings for each sample combination. These tests were performed for the entire user group as well as for the subset of non-novice users, the results of which are shown graphically in Figure 4.18. The experienced user subgroup was chosen to isolate the effects of first-time users, whose ratings varied substantially between trials, as will be discussed below.

While the majority of sample rating pairs showed little correlation for either group, there were some noticeable exceptions. For the event-based samples, we see that the strength of the underlying proportional controller did not significantly affect realism ratings, with p-values ranging from 0.21 to 0.82. This trend can be seen by the clustered elements near

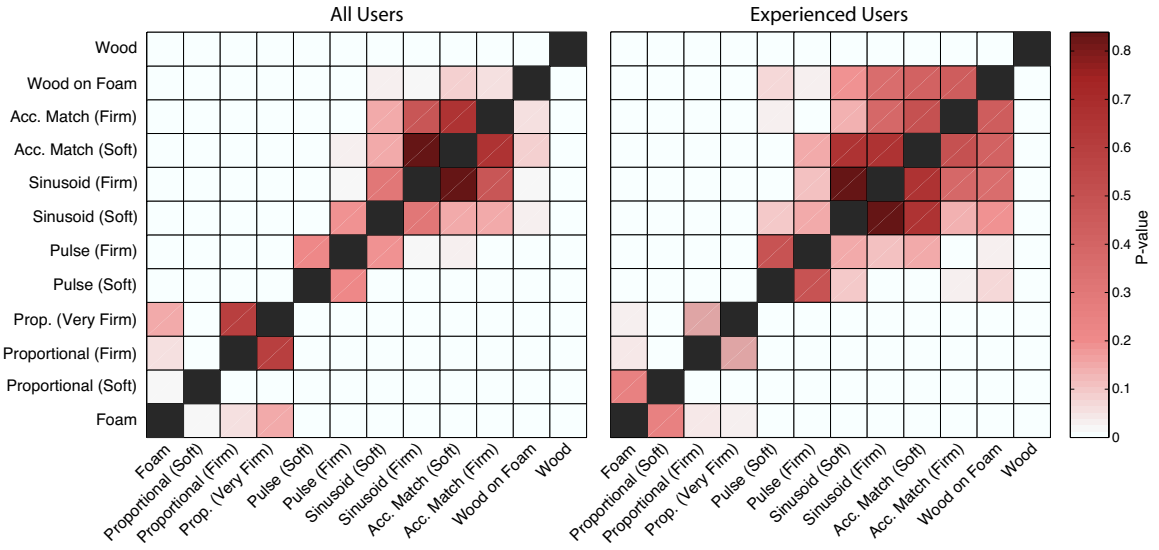

Figure 4.18: P-values from t-tests on the average realism rating given by each subject for all sample pairs: the shade of each square shows the probability that the ratings given to the two intersecting samples stem from indistinguishable populations.

the diagonal in both charts of Figure 4.18. Such a finding highlights the salient role that high-frequency transients play in haptic perception of contacts.

In comparing the results of the entire group with the subset of experienced users, we see that the experienced group ratings for the event-based samples and the wood-on-foam sample are statistically very similar (p-values ranging from 0.20 for the soft sinusoid to 0.47 for the firm acceleration matching). These results echo closely those found in our preliminary study, which included a higher percentage of expert users [78]. This finding is particularly interesting because all of the event-based virtual samples were constructed to match the experience of tapping on the wood-on-foam sample (see Section 4.5.2), demonstrating the efficacy of the event-based paradigm for mimicking real contact transients.

While many factors likely contributed to the different ratings provided by the experienced group and the novice users, the testing experience may have been somewhat overwhelming to the novices. Many first-time users appeared to have difficulty identifying the subtleties of the various samples, and a few admitted that their ratings were arbitrary at times. In contrast, those with some level of experience seemed to be more capable of identifying the differences and gave qualitative feedback that was much more consistent.

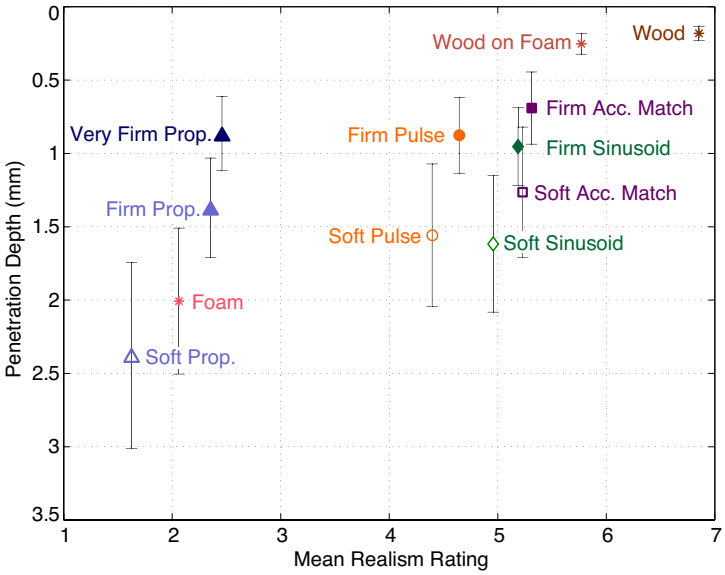

Figure 4.19: Mean penetration depth versus realism rating for each sample: error bars indicate one standard deviation from the mean.

Alternative hypotheses include speculation that naïve users are more attuned to the characteristics of real interactions, as well as a conjecture that experienced users might be more accepting of virtual rendering methods due to their familiarity.

Rating Criteria

When asked to name their rating criteria, subjects listed several salient metrics. First among these was whether the stylus came to a sudden stop after contact. In our preliminary study [78], we hypothesized that the foam and the simple proportional controllers were rated most poorly because they cannot quickly cancel the user's incoming momentum. This study tracked penetration depth during every tapping event to test this hypothesis. Figure 4.19 shows the mean penetration depth for each sample, plotted against the sample's average realism rating. While the wood and wood-on-foam sample penetration depths were less than 0.25 millimeters, the foam and soft proportional controller allowed penetrations of over 2 millimeters. Interestingly, we discovered that the average penetration depth for the firm and very firm proportional samples were actually less than many of the event-based samples, but their ratings were consistently much lower. This finding alone demonstrates that there is clearly more to impacts than can be represented with simple position-based control.

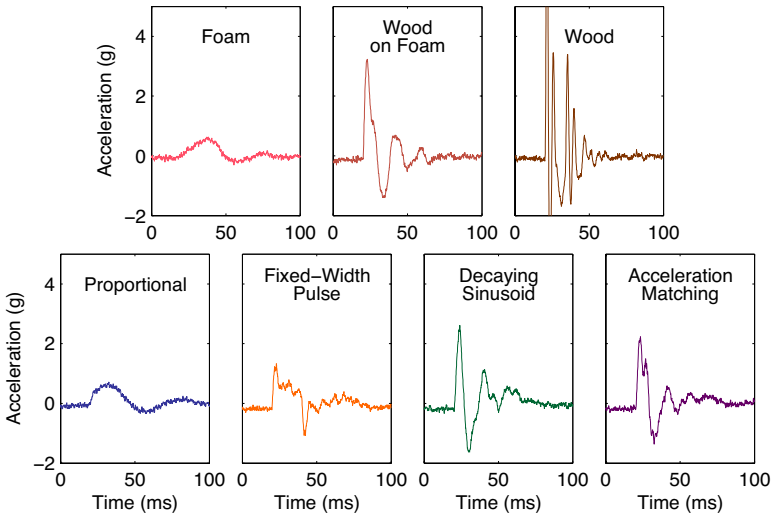

Figure 4.20: Typical accelerations for contact with samples at an incoming velocity of 0.11 meters per second.

The second most commonly mentioned criterion was the presence of a high-frequency transient at the moment of contact. Re-examining the results shown in Figure 4.19, we see that all of the event-based samples were rated more highly than the proportional samples, a result also seen clearly in Figure 4.17. Directly comparing the penetration depths for the event-based samples and their underlying proportional controllers, we see that the inclusion of the transient roughly decreases the penetration depth by a factor of two. The transient cancels a significant portion of the user's incoming momentum, decreasing penetration and improving the realism of the virtual surface.

We can also examine the accelerations produced by tapping on the real and virtual surfaces, as shown in Figure 4.20. Firm and soft underlying proportional control produce nearly identical acceleration histories and realism ratings, so only the firm samples are illustrated. The three event-based virtual samples produce high-frequency accelerations that are similar to those seen for the wood on foam and wood objects. Of these three, the decaying sinusoid and acceleration-matched transients most closely resemble the real signals, which we hypothesize contributes to their high realism ratings.

As additional criteria, users commented that they gave lower ratings to samples that felt unnatural. Two subjects noticed a high-frequency buzzing that detracted from sample realism. I believe this sensation stems from encoder discretization in the proportional controller that holds the stylus vertical during all tests, but it may also be a result of

encoder discretization in the firm and very firm virtual objects. Also, a few users reported feeling an occasional double impact, which can occur when the stylus is held very loosely. I believe that measuring and compensating for changing hand impedance could further improve the realism of event-based transient display, avoiding excessive force application and double event triggering.

These human subject evaluations confirm that high-frequency transients are vital to achieving realism, proving substantially more important than penetration distance or object stiffness. Outputting either manually tuned decaying sinusoids or analytically computed acceleration-matched signals provides the highest level of realism; I believe that these methods succeed because they create high-frequency stylus accelerations that closely match those experienced during real contact. The HFAM transients are generated automatically and can be applied to any surface interaction encountered, adding to the versatility of this approach. Experienced haptic users cannot distinguish such virtual event-based renditions from a piece of wood mounted on a similarly soft substrate. Users also judge classic proportional feedback to be equivalent to real foam, reiterating the softness of traditional haptic displays.

4.6 Summary

Humans rely on information-laden high-frequency accelerations in addition to quasi-static forces when interacting with objects via a handheld tool. Telerobotic and virtual environment systems have traditionally struggled to portray such contact transients due to closed-loop bandwidth and stability limitations, leaving remote objects feeling soft and undefined. High-frequency acceleration matching builds on prior efforts in vibration feedback to enable a haptic interface's standard motors to create fingertip contact accelerations that feel like real interactions.

The fundamental insight to the HFAM approach is that the internal dynamics of the user-interface system color any high-frequency feedback displayed with the motors. These dynamics can be characterized offline via the methods of Chapter 3, estimating the relationship between commanded motor current and handle acceleration while a user holds the device. The resulting model can then be inverted to allow computation of the current command necessary to cause a certain high-frequency acceleration at the handle of the interface. The technique of acceleration matching can be applied to haptic interaction with both remote and virtual environments.

In teleoperation, a high-bandwidth sensor measures accelerations at the slave's tip, and the real-time controller recreates these important signals at the master handle in combination with position-error-based feedback of quasi-static forces. This hybrid signal closely corresponds to the asymmetry of human sensing capabilities, instilling telerobotics with a more realistic sense of remote touch. One interesting feature of this algorithm may arise when it is extended to systems in which the master and slave have substantially different masses. With a large slave, the accelerations caused by environmental contact will be lower in frequency than those expected by the user. With a small slave, such as a micro-electro mechanical system, contact accelerations will contain very high frequencies that may be beyond the user's one kilohertz sensory bandwidth. In these cases, the recorded slave acceleration may need to be frequency-shifted or otherwise adjusted before playback to achieve a natural feel for the user.

In virtual environments, high-frequency acceleration matching pre-computes impact transients from recorded contact acceleration traces, yielding a library of event-based feedback. Standard position feedback is then augmented by displaying these transients open-loop when a contact event is triggered. The current command transients account for the internal electrical, mechanical, and biomechanical dynamics of the haptic interface via model inversion. These techniques enable a haptic interface designer to output specified high-frequency accelerations at contact: such tools may prove useful for understanding the relationship between high-frequency fingertip acceleration and the haptic impression of contact. For example, I have found in both teleoperation and virtual environments that care should be taken to maintain feedback levels that do not overwhelm the user's incoming momentum; scaling the magnitude of virtual transients down by ten to 20 percent keeps the surface feeling passive and natural without sacrificing a great deal of realism. The relationship between portrayed acceleration and perceived sensation deserves further investigation.

This research found that adding high-frequency accelerations to simulated haptic interactions makes them feel far more real than the contacts portrayed by traditional feedback methodologies. The realism of the HFAM approach to virtual tapping feedback was evaluated by a group of users in a blind study. Users gave the acceleration-matched virtual samples a median rating of about 5.5 out of 7.0, compared with 2.0 for proportional feedback and 6.0 for the sample of real wood on foam. Encouraged by the results of this work, I firmly believe the paradigm of high-frequency acceleration matching has the potential to significantly improve haptic display. This approach particularly improves the rendering of hard

contact, which is an important challenge in haptic simulation, and it may provide valuable extensions in the domain of texture as well. Just as finite element methods and extensive off-line computation have improved the haptic fidelity of deformable objects, high-frequency acceleration matching enables accurate rendering of the dynamics that characterize rigid contact.

Chapter 5

Canceling Induced Master Motion

Teleoperation systems have persistently struggled to provide users with realistic force feedback; high-frequency contact transients convey important information about the remote environment but must typically be attenuated to avoid the contact instability they incite. This undesirable behavior can be traced to induced master motion, movement of the master device that is caused by haptic feedback rather than user intention, as defined in Section 5.2. Such motion is interpreted as a position command to the slave, closing an internal control loop that is unstable under high gain. The resulting behavior of contact instability is analyzed from a momentum transfer standpoint in Section 5.3. Instabilities of this type can occur in any teleoperation control architecture that uses the master device to apply large, high-frequency haptic feedback while measuring its position for the slave command: the classic example is a position-force controller, which is formally described in Section 5.1, and HFAM-augmented position-position control may engender similar behaviors on some systems. This phenomenon may also become relevant to virtual environments as their haptic feedback algorithms incorporate stronger event-based transients like those afforded by high-frequency acceleration matching.

After prior efforts to stabilize high-frequency haptic feedback are discussed in Section 5.4, Section 5.5 presents a new approach for achieving stable, high-gain force reflection via model-based cancellation. I call this strategy canceling induced master motion or CIMM; it was first proposed in [80], and it is fully detailed in [83]. This chapter focuses on applying cancellation to position-force teleoperation, but its methods could also be applied to haptic interfaces used in other types of teleoperation or in virtual environment simulations. Requirements for the model of the induced motion dynamics are described in

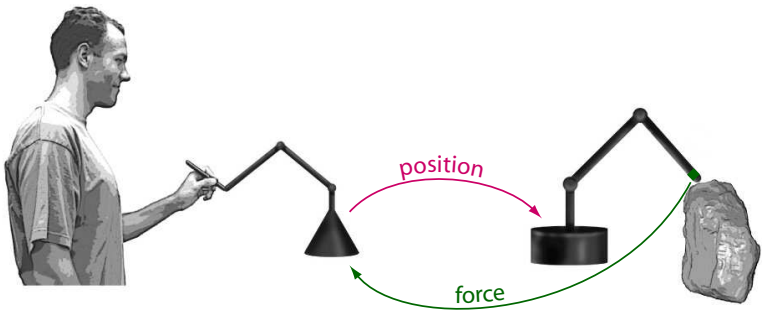

Figure 5.1: Position-force control for teleoperation.

Section 5.6, building on the characterization tools developed in Chapter 3. Cancellation control is tested in Section 5.7 on a one-degree-of-freedom teleoperation system, using a sixth-order nonlinear model of the user-master dynamics. As summarized in Section 5.8, canceling high-frequency induced master motion during teleoperation improves the stability of impacts, allowing significantly higher force reflection levels and a more authentic user experience.

5.1 Position-Force Control

Among the many controllers used in telerobotics, the position-force architecture seeks to provide accurate feedback by explicitly measuring the force of contact between the slave and the environment. As illustrated in Figure 5.1, the slave is commanded to follow the measured position of the master mechanism. Forces sensed at the slave's end effector are simultaneously displayed via the master's motors, transmitted to the user's hand via the structure of the device. One fascinating aspect of this architecture is as follows: throughout an interaction, the master mechanism must perform the dual tasks of position measurement and force display. The dynamics of the master device couple these two functions together, allowing the controller's force output to affect its position input.

Like a sound system whose microphone is too close to the speakers, telerobotic systems let high-frequency force feedback induce motion of the master device that is then treated as a position command. During environment contact, this induced motion can drive the entire system unstable with high-frequency vibrations, similar to the screeching of a badly configured sound system. As with speaker volume, the force feedback gain must usually be lowered to achieve stability, reducing haptic cues to the user and leaving interactions feeling

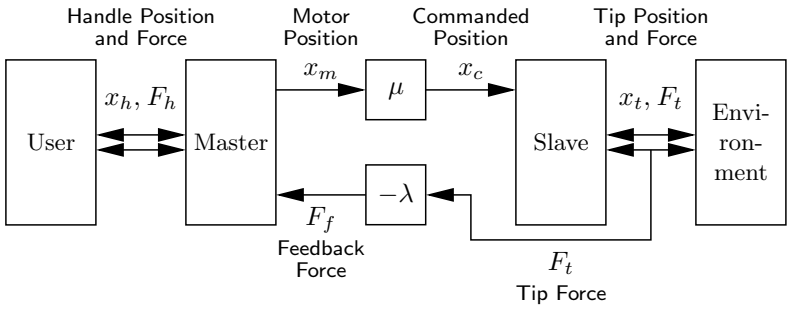

Figure 5.2: Block diagram of position-force control.

soft and ill-defined. This interesting, undesirable behavior can be understood by examining the dynamics of the system more carefully.

5.1.1 Architecture Definition

The position-force architecture is drawn as a block diagram in Figure 5.2, showing user, master, controller, slave, and environment. During teleoperation, the user holds the handle of the master mechanism and moves it around. Assuming a causality for this interface is not required for this work, so the interaction is depicted with two double-headed arrows. The system monitors the position of the master mechanism's motors, x_m, and commands the slave robot to move accordingly via x_c. The slave controller attempts to make the slave track this motion command, typically via proportional, integral, and/or derivative control. We draw the interaction between the tip of the slave arm and the environment as bidirectional in position and force, again to avoid an unnecessary assumption of causality. As the slave moves, the system captures the effects of environmental contact by measuring the force exerted by the tip, F_t. A scaled version of this tip force is continuously displayed via the motors on the master device as F_f to allow the user to feel the interaction.

5.1.2 Gain Selection

The forward position scaling ratio, μ, and the force feedback gain, $-\lambda$, can be used to adjust the interaction between the two sites [22], following the naming convention of Daniel and McAree [28]. Typically, μ is chosen to scale comfortable human hand movements to the workspace of the slave robot, allowing the user to perform the chosen task as though it

were a familiar manual manipulation. Common values of μ include 0.004 for nanomanipulation [138], about 0.33 for minimally invasive surgery [125], and 1.0 for nuclear decommissioning [28].

Once μ is selected, an appropriate value for λ must be found. In the same way that human motion is scaled to match the task, contact forces should be scaled up or down to match human sensory capabilities [49]. Human users want to comfortably feel the slave's interaction with the environment. Unfortunately, stability concerns usually limit the force reflection ratio to values below this desired level [88]. In practice, researchers begin by setting $\lambda = 0$ to verify that the slave can make stable contact with the environment in the absence of force feedback. To haptically augment the user's experience, λ is then gradually increased until the slave begins to exhibit unnatural vibrations when touching hard objects. At this relatively low value of λ, soft objects are generally difficult for the user to detect, and hard objects feel like foam or rubber rather than wood or metal. Systems that can support higher values of λ for a given μ will provide stronger haptic cues to the user, allowing more natural interactions and greater sensitivity during delicate operations such as microsurgery.

5.2 Induced Master Motion Pathway

High-gain position-force teleoperation is plagued by instability; designers commonly encounter a maximum force feedback gain, λ, above which their system cannot make stable contact with hard objects in the environment. This instability occurs because a teleoperator is a closed-loop control system whose internal dynamics change with the gains μ and λ, as well as with the natural responses of the user, master, slave, and environment. The signal loop in question can be seen in Figure 5.2, passing from master motor position to commanded position, slave tip force, feedback force, and back to master motor position through a series of connecting elements. The first three elements of this loop are required for teleoperation: μ conveys the user's intended motion, the slave interacts with the environment to accomplish the task, and $-\lambda$ reflects the contact force to the master for the user to feel. The element that closes this potentially unstable loop is the user-master system, a connection that is commonly acknowledged but to date has not been highlighted during analyses of closed-loop performance.

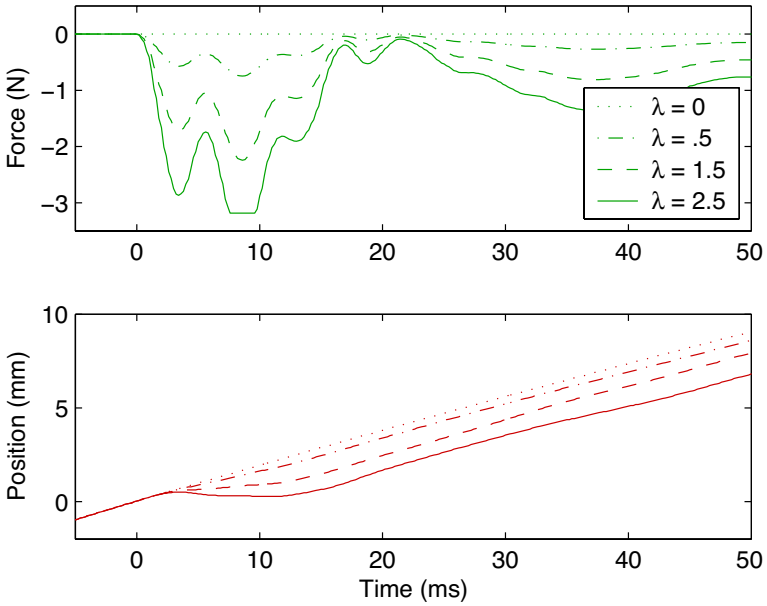

Figure 5.3: Display of a scaled, pre-recorded force profile to a user executing a constant motion with a one-dof master.

5.2.1 Voluntary and Induced Master Movement

The master mechanism plays an important bidirectional role in a teleoperation system, connecting the user to the controller via the series of dynamic elements illustrated in Figure 3.1. The user's hand acts on the master at the handle while the controller applies the feedback force, F_f, at the motor. These two interfaces jointly determine the movement of the device and thus the position of the master motor, x_m. In the absence of force feedback, the human is the only active influence on the master mechanism, and x_m is a good measurement of his or her smooth, low-frequency hand motion. When the controller applies additional forces at the motor for the user to feel, these forces interact with the dynamics of the user-master system and cause movement that the human does not intend.

To illustrate this phenomenon, Figure 5.3 shows the first fifty milliseconds of measured motor motion during open-loop force display on a one-dof master. The user moved the handle of the device at a constant velocity, and a scaled, pre-recorded force profile was displayed when the motor crossed a position threshold. The dynamic coupling between force feedback and motor position is visible as sudden deviations from the user's smooth, $\lambda = 0$ path, happening far faster than the human's cognitive reaction to the applied force,

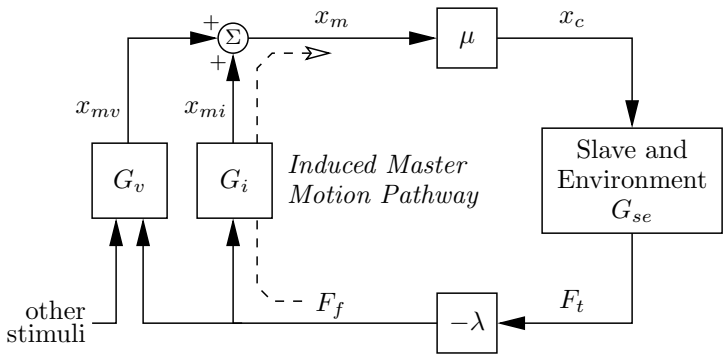

Figure 5.4: Dual motion pathways of the master motor.

which takes approximately 140 milliseconds [92]. Higher levels of force reflection provide stronger haptic cues to the user but also cause more significant divergence from the user's smooth, intended path. I call these deviations *induced master motion* because they stem from the force feedback signal rather than the user's volition.

5.2.2 Superposition Assumption

These dual sources of master movement carry through to teleoperation, where the measured position of the master, x_m, is determined both by the user's active influence on the handle and by the system's passive response to the feedback forces applied at the motor [28, 39]. Following this observation, I can redraw the standard position-force teleoperator as shown in Figure 5.4, separating voluntary and induced master motion into parallel pathways and also lumping the slave and environment into a single entity. I approximate the master motion signal via superposition as

$$x_m(t) = x_{mv}(t) + x_{mi}(t), \tag{5.1}$$

where x_{mv} is the user's voluntary contribution, x_{mi} is the component induced by force feedback, and t is time. Furthermore, I define the connection from force feedback to unintentional motor movement as

$$x_{mi}(t) = G_i\left(F_f(t)\right), \tag{5.2}$$

where G_i represents the induced motion dynamics of the user-master system and may be nonlinear and/or time varying. This function depends on electrical and mechanical characteristics of the device and on the instantaneous biomechanical properties of the user's

hand and arm; as such, it can be modeled using specialized system identification techniques like those developed in Chapter 3.

In the parallel pathway, the function G_v represents the user's voluntary (active) response to a wide array of received stimuli, including auditory, visual, and haptic feedback. I do not attempt to model this highly cognitive process and instead treat x_{mv} as an exogenous input, noting the human ability to stably interact with and even stabilize typical dynamic systems [61].

The underlying assumption of superposition between voluntary and induced master motion is reasonable because the two signals differ in magnitude and frequency content. The user's voluntary hand movements are on the order of centimeters and consist of frequencies below five hertz, which are well transmitted from the handle to the motor by the impedance-type devices with which I am concerned. On the other hand, the induced master motion in these stiff systems is typically on the order of millimeters and contains a broad range of frequencies which cannot naturally appear in x_{mv}. I thus treat x_{mi} as a disturbance to the user's voluntary path and investigate its implications for system performance. It is worthwhile to note that this partitioning is consistent with the commonly employed user-master model of a mass, spring, and damper connected to the user's slowly changing desired position [28], though more detailed models are required to accurately capture the broad-frequency behavior of induced master motion.

5.3 Contact Instability

During teleoperation, the pathway of induced master motion interacts with slave and environment dynamics to form an internal controller loop that includes both μ and λ, as can be seen in Figure 5.4. From this diagram, I can write out the equation for the master position, x_m, proceeding once around the loop:

$$x_m(t) = x_{mv}(t) + G_i\left(-\lambda G_{se}\left(\mu x_m(t)\right)\right). \tag{5.3}$$

The physical interpretation of this signal loop is that the natural dynamics of the user-master system, G_i, allow the controller's force feedback signal to induce motion of the master motor that is not intended by the human. This induced motion is interpreted as a slave movement command, which causes tip motion and additional force feedback from the environment. Ideally, the slave's motion command would depend only on the user's

voluntary motion, x_{mv}, and not on the system's closed-loop dynamics, which can distort or destabilize the interaction.

If we assume that G_i and G_{se} are linear time-invariant (LTI) systems, we can formulate the system transfer function from (5.3) as follows:

$$\frac{x_m(s)}{x_{mv}(s)} = \frac{1}{1 + \mu\lambda G_i(s) G_{se}(s)}, \tag{5.4}$$

where s is the Laplace operator. Although they are not valid for most real systems, and especially not for a slave that is making and breaking contact with the environment, these assumptions are widely applied in the literature and can provide useful insights on the behavior of the system. Specifically, the denominator of (5.4) is the system's characteristic equation, the roots of which are the closed-loop poles of the teleoperator. For the system to be stable, all of the poles must be in the left half of the s-plane, a criterion that depends on μ, λ, $G_i(s)$, and $G_{se}(s)$. For a position-force controlled system with a given $G_i(s)$ and $G_{se}(s)$, stability is determined by the product of μ and λ. Once the position scale is chosen, there generally exists a force feedback gain above which the slave cannot maintain stable contact with stiff objects.

5.3.1 Simple Stability Limit

Daniel and McAree analyze this limit intuitively through conservation of momentum as well as mathematically via root-locus techniques [28]. They model both master and slave as pure masses, eliminating the distinction between master handle and motor as well as between slave motor and tip. They begin by considering a situation in which the two devices are tracking perfectly, with the slave moving at a constant velocity $v_s = v$ and the master moving at $v_m = v/\mu$. An impact with the environment will change the slave's momentum by ΔP. If this impulse is reflected perfectly with a force-feedback gain of λ, the master's momentum will change by $\lambda\Delta P$. Letting m_m and m_s stand for the effective masses of the master and the slave, and letting δv_m and δv_s represent the respective changes in velocity that result from contact, I can write

$$\Delta P = m_s \delta v_s = \frac{m_m \delta v_m}{\lambda}. \tag{5.5}$$

In my architecture, δv_m is the master motion induced by the contact and can be written as

$$\delta v_m = \lambda \frac{m_s}{m_m} \delta v_s. \tag{5.6}$$

Even in this simplified treatment of system dynamics, induced master motion can be seen to increase with higher force feedback gains and to depend strongly on the dynamics of the user-master system, which are represented here by m_m.

After such an impact, the slave will be able to maintain contact with the environment only if its command does not pull it away from the contacted object at a rate faster than its natural rebound. This relationship requires the following inequality between post-impact slave and command velocities:

$$\begin{aligned} v_s' &< v_c' \\ v + \delta v_s &< \mu(v/\mu + \delta v_m), \end{aligned} \tag{5.7}$$

taking v to be positive and δv_s and δv_m to be negative. By subtracting v from both sides of (5.7), substituting (5.6), and dividing by the negative value δv_s, we obtain the following requirement for contact stability on an otherwise uncompensated system:

$$\mu\lambda < \frac{m_m}{m_s}. \tag{5.8}$$

Though it presents a significantly simplified view of contact dynamics, this naïve analysis supports my identification of induced master motion as the culprit of contact instability during teleoperation, as well as the stabilizing effect of decreasing the force feedback gain. Daniel and McAree derive a more detailed criterion for their simple system model via root locus methods, again finding an upper limit for $\mu\lambda$ that depends on the dynamics of the teleoperator.

5.3.2 Master Dynamics and System Stability

Although a detailed derivation is beyond the scope of this dissertation, I assert that the threshold of stability in the general case is determined by the dynamic interaction between the induced master motion pathway, G_i, and the slave-environment system, G_{se}. Widely observed in practice, the slave enters a limit cycle at high values of λ, repeatedly making and breaking contact with the environment. In addition, moderate values of λ for a given μ

prolong the contact transient, distorting the natural response of the environment. Lowering the force reflection ratio re-establishes functional operation but limits the user's ability to discern environment properties. Ideally, user preference and task requirements would guide selection of λ, rather than closed-loop stability.

Little research acknowledges the pivotal role that the passive dynamics of the user-master system play in this limitation. As discussed in Section 3.2.2, most analyses model the master as a pure mass and the user as a spring and damper even though these dynamics are truly more complex, especially at high frequency. Identifying the actual relationship between F_f and x_{mi} can contribute to an understanding of overall system behavior, including a derivation of a system's actual stability threshold. Master devices can be evaluated according to the G_i they provide, as devices that permit less induced master motion will support the higher force feedback levels that are commonly desired. Furthermore, the electrical and mechanical components of existing master systems can be modified to adjust G_i and increase the force-feedback gain that they will stably support. In summary, induced master motion compromises the stability of a telerobot by allowing the force feedback signal to influence the slave's commanded position, creating an internal control loop that is unstable under high gain.

5.4 Prior Work in Stabilizing Haptic Feedback

Recognizing that teleoperative contact instability occurs when the force-feedback gain λ is too large for a given position scale μ, many researchers have proposed strategies for handling this trade-off. As mentioned above, the simplest approach lowers the force feedback gain until contact with the stiffest environment elements becomes stable. Such a choice reduces the intensity of the user's force cues and often leaves lightly damped oscillations that can distract from the interaction. Contact forces can be portrayed via auditory or visual feedback instead, but such an arrangement sacrifices the natural connection between hand forces and hand motions.

Another basic tactic requires the user to hold the master mechanism with a firmer grasp or even with two hands. Increasing the user's impedance in such a way attenuates handle motion and generally decreases induced motor motion, lowering the magnitude of G_i. Although this adjustment can be voluntary, the movement that results is still termed induced master motion because the user cannot prevent it from occurring. A firmer user grasp

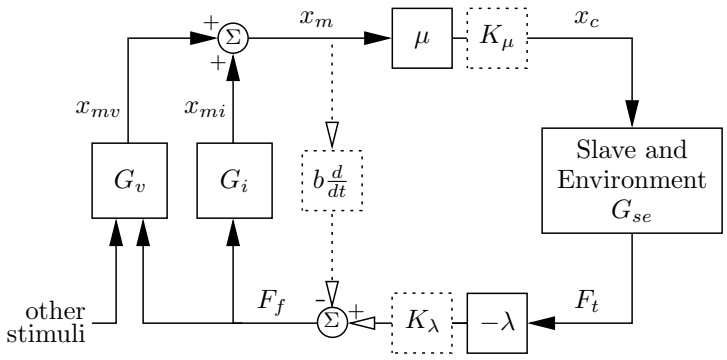

Figure 5.5: Position-force control with three compensation options: local derivative feedback on master position via $b\frac{d}{dt}$, position command filter K_μ, and force feedback filter K_λ.

can quell some marginally unstable oscillations, but it is not a solution to the underlying problem.

More sophisticated methods are required to overcome the instability generated by position-force control. Many proposed strategies aim to accomplish this goal by shaping the signals of the system's internal feedback loop. Three standard stabilization approaches are illustrated in Figure 5.5, including local derivative feedback on the master, a position-command filter, and a feedback-force filter. These modifications can be used alone or in tandem to adjust the behavior of a troublesome system, as discussed below.

5.4.1 Master Damping

The first illustrated approach for suppressing induced master movement is to add damping to the master manipulator via local feedback [13, 55]. Augmenting the block diagram with negative derivative feedback from x_m to F_f changes the system's loop relationship to

$$x_m(t) = x_{mv}(t) + G_i\left(-b\,\dot{x}_m(t) - \lambda\, G_{se}\left(\mu\, x_m(t)\right)\right). \tag{5.9}$$

The added term punishes high motor velocities via an opposing control force, which is added to the force feedback signal before being applied to the motor. Although it can stabilize contact, this additional damping is also apparent to the user and makes the system feel slow and unresponsive, even when the slave is moving in free space [28]. Furthermore, this strategy requires a clean, accurate measurement of velocity, which is seldom available in systems that rely on numerical differentiation of discrete position signals.

5.4.2 Position Command Filter

Other researchers have explored the possibility of putting a compensator in the slave's commanded position via K_μ [28, 39]. The loop equation then becomes

$$x_m(t) = x_{mv}(t) + G_i\left(-\lambda\, G_{se}\left(K_\mu(\mu\, x_m(t))\right)\right). \tag{5.10}$$

When K_μ is a low-pass or notch filter, it attenuates the position command's destabilizing frequency content, but it can also prevent the slave from tracking quick movements by the user. In the general case, the compensator must be carefully selected to avoid adding too much lag at the crossover frequency, which would compromise the system's stability margins.

5.4.3 Force Feedback Filter

The third main stabilization strategy involves low-pass or notch filtering the force feedback signal via K_λ, removing mid- to high-frequency content before display on the master mechanism [28, 54]. This compensator yields

$$x_m(t) = x_{mv}(t) + G_i\left(K_\lambda\left(-\lambda\, G_{se}\left(\mu\, x_m(t)\right)\right)\right). \tag{5.11}$$

Filtering the force feedback can stabilize a system, but it also prevents the user from feeling high-frequency signals, resulting in interactions that feel soft and undefined. Such a choice enables stability under higher force reflection ratios, but it compromises the information content of the feedback signal, especially when initiating contact with a stiff environment.

5.4.4 Vibrotactile Feedback

The feel of a system with filtered force feedback can be improved by using a separate actuator to display the signals removed by K_λ, a technique often described as combined vibrotactile and force feedback [76]. This frequency-domain separation increases the information available to the user while maintaining system stability. It should be noted that the vibrations from the additional actuator must be carefully isolated from the forward position command so that they do not enter the closed loop and incite contact instability. This strategy of combining actuators effectively bypasses the pathway of high-frequency induced master motion, staving off instability while allowing the user to feel a fuller spectrum of environment feedback.

5.4.5 Alternative Methods

Beyond these four standard approaches, some unique alternatives have also been considered. The controller described by McAree and Daniel [102] anticipates collisions between the slave and the environment using external optical sensors. When contact is imminent, the controller actively reduces the slave's velocity to minimize the forces generated, effectively adjusting $G_s(s)$ in real time. This approach requires less force-feedback filtering than an uncompensated system does, allowing the user to feel more mid-frequency signals; however, this adjusted slave behavior can be disorienting to some users [102].

Another strategy involves continually predicting the user's intended position command via a model of the human rather than using the presently measured master position. Explored options range from polynomial and spline extrapolation [122] to a full human arm model with measured neural inputs [121]. Most such efforts have difficulty achieving high-bandwidth position tracking, which is imperative for transparency, but the focus on human intention rather than measured master position is intriguing. Clearly, the dynamic interaction between human, master, slave, and environment merits further exploration.

5.5 Cancellation Approach

As an alternative to standard loop-shaping strategies and additional actuators, we can use the perspective of induced master motion to break the internal loop of the controller via model-based high-frequency cancellation. First proposed in [80] and developed further in [83], the cancellation approach centers on understanding the passive response of the user-master system during a telerobotic interaction to improve high-gain system stability. Here, we will examine cancellation through simple momentum-based arguments, the more general architecture of parallel master motion pathways, and the commonly used four-channel architecture for teleoperation.

Although new in telerobotics, canceling induced master motion bears a resemblance to several existing model-based control strategies. Some researchers attempt to hide the dynamics of the master from the user to make it feel like a weightless bar rather than a physical device [91,156]. Typically modeling the master as a simple mass with viscous damping, they apply additional feedback forces to cancel the inertial and frictional forces estimated from position measurements. Cancellation also bears a resemblance to command input shaping, which has been applied to the control of contact transitions in robot manipulators [65],

among other applications: the system's input is convolved with its response to suppress easily excited oscillations at contact events, a process that is somewhat analogous to my approach of removing the master's response to force feedback from the slave's position command. Another application of model-based cancellation is biodynamic feed-through, where accelerations of a mobile vehicle apply forces to the body of the operator and therefore to the on-board control joystick. With a simple, second-order model of these dynamics, one can avoid instability by applying equal and opposite torques to the joystick [44, 137]. I extend these force-canceling ideas to cancellation of induced master motion, where I want the operator to feel the applied feedback without allowing the forces to affect the slave's position command.

5.5.1 Simple Stability

If we return to the naïve stability analysis of Section 5.3.1 and look again at (5.7), we see that the velocity of the slave command after impact, v_c', is affected by δv_m, the natural response of the master mechanism to the feedback force. When the slave command is formed directly from the master position, as is typical, the master's induced movement sets the stability limitation derived in (5.8). With knowledge of the master's dynamics, however, we can estimate the master movement that will result from the impact, $\hat{\delta v}_m$, and remove it from the slave command, as

$$v_c' = v + \mu\,\delta v_m - \mu\,\hat{\delta v}_m \approx v. \tag{5.12}$$

Accurately anticipating and canceling induced master movement provides a post-impact slave command that depends only on v, the pre-impact slave velocity, which embodies the user's voluntary movement. The inequality necessary for stability,

$$\begin{aligned} v_s' &< v_c' \\ v + \delta v_s &< v, \end{aligned} \tag{5.13}$$

then always holds, as v is positive and δv_s is negative, and the system is theoretically stable for all gains.

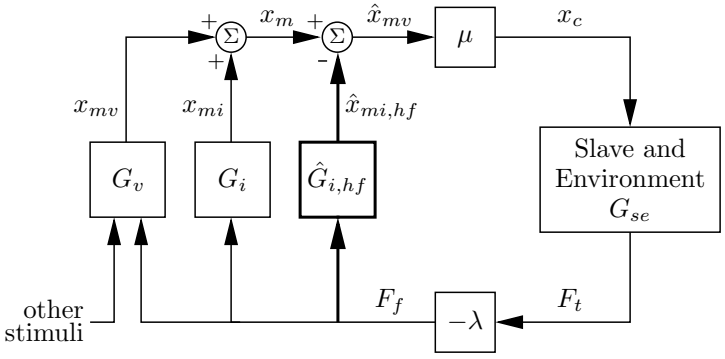

Figure 5.6: Canceling high-frequency induced master motion.

5.5.2 Induced Motion Pathway

Cancellation is illustrated for our more general architecture in Fig. 5.6; as described in Sec. 5.2, force feedback induces deviations from the user's intended path through the element G_i. Deviations that are within the user's motion control bandwidth of eight hertz, such as the deflection of the device and hand under a constant contact force, are naturally compensated for by the user via changes in x_{mv}. But the user cannot actively compensate for induced master motion above about eight hertz, leaving the system vulnerable to instability. In this new control method, I model the high-frequency relationship between F_f and x_{mi} as $\hat{G}_{i,hf}$ and use it to cancel induced master motion via real-time simulation. The model's response, $\hat{x}_{mi,hf}$, is subtracted from the measured master position, x_m, to provide an estimate of the user's intended path, $\hat{x}_{mv}$. This estimate is multiplied by μ to become the slave's position command, x_c, providing a signal that is ideally free from both low-frequency and high-frequency feedback artifacts.

With cancellation in place, the system's loop relationship becomes

$$\begin{aligned}\hat{x}_{mv}(t) = x_{mv}(t) \; &+ \; G_i\left(-\lambda\, G_{se}\left(\mu\, \hat{x}_{mv}(t)\right)\right) \\ &- \; \hat{G}_{i,hf}\left(-\lambda\, G_{se}\left(\mu\, \hat{x}_{mv}(t)\right)\right), \end{aligned} \tag{5.14}$$

with a transfer function of

$$\frac{\hat{x}_{mv}(s)}{x_{mv}(s)} = \frac{1}{1 + \mu\lambda G_i(s) G_{se}(s) - \mu\lambda \hat{G}_{i,hf}(s) G_{se}(s)} \tag{5.15}$$

under LTI assumptions. If the $\hat{G}_{i,hf}$ model perfectly captures the high-frequency dynamics

of induced master motion, the system's characteristic equation will be unity, indicating the absence of closed-loop dynamics. The connection from feedback force to slave motion command will be severed, and the system will no longer exhibit contact instability. Although perfect cancellation is unattainable, this control methodology can achieve substantial performance gains in real systems when the model accurately captures the high-frequency dynamics in question.

The key step to this approach is obtaining a good model, $\hat{G}_{i,hf}$, which will be the subject of Section 5.6. The model must capture the high-frequency response of the master motor to force feedback as a user holds the device's handle. The more closely the model approximates real system behavior, the more attenuated the connection from F_f to x_c becomes, stabilizing the system for higher force-feedback gains. A decent approximation of the system's high-frequency response under the conditions that incite instability will reduce the controller's loop gain and stabilize contact for moderate values of λ. The user will be able to feel strong feedback forces, and the slave command will closely track x_{mv}, the path that the master motor would have taken in the absence of force feedback.

5.5.3 Four-Channel Architecture

Cancellation of induced master motion can also be understood by formulating it in teleoperation's four-channel architecture, which was first presented by Lawrence [88] and expanded by Hashtrudi-Zaad and Salcudean [57]. Such formulation requires three assumptions beyond those already in place. First, linear models are used for the human, master, slave, and environment. Second, an impedance causality is assigned to both the human and environment. And lastly, human forces and motor forces are collocated, acting on the master with the same transfer function. Nevertheless, this architecture provides a useful perspective on the cancellation approach.

A position-force controller with cancellation is illustrated in Figure 5.7 using the common nomenclature for the various subsystem blocks. Elements that are not typically used in position-force control are not illustrated, i.e. $C_3 = C_4 = C_6 = C_m = 0$. The values for the

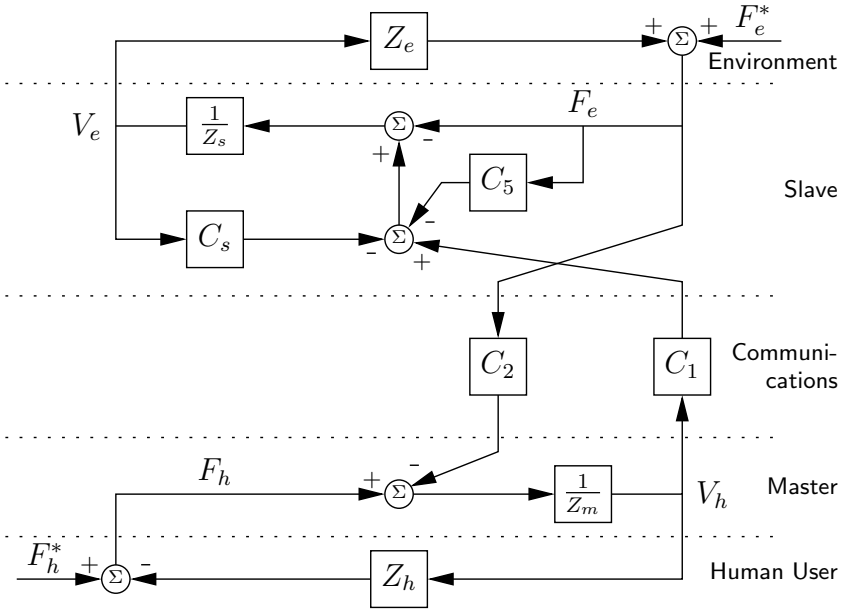

Figure 5.7: Four-channel treatment of position-force control with cancellation of induced master motion.

remaining elements are

$$C_1 = \mu(b_s + k_s/s) \tag{5.16}$$

$$C_2 = \lambda \tag{5.17}$$

$$C_s = b_s + k_s/s \tag{5.18}$$

$$C_5 = -\lambda\mu\frac{(b_s + k_s/s)}{\hat{Z}_m + \hat{Z}_h} \tag{5.19}$$

where b_s and k_s are slave controller gains and s is the Laplace operator. Cancellation of induced master motion is embodied in the C_5 element, which is usually regarded as a local force control loop on the slave. It works to prevent the environmental contact force, here F_e, from contributing to the slave's control force through the passive dynamics of the master and user. If the master and human models ($\hat{Z}_m$ and $\hat{Z}_h$) are accurate, the net force applied to the slave via C_1, C_s, and C_5 will represent only the user's voluntary influence on the master, F_h^*. The slave command will be free from any feedback artifacts, and the system's troublesome closed loop will be broken. Theoretically, such a system would support arbitrary gain choices, though in practice model fidelity becomes the limiting factor to performance.

5.6 Modeling Induced Master Motion

To achieve the benefits of the cancellation approach, one must obtain a good model of the relationship between force feedback and high-frequency induced motor movement; denoted $\hat{G}_{i,hf}$, this model of the user-master system need not be linear or time-invariant. As illustrated in Fig. 3.1, the target system extends from the master's amplifier and motor through its cables, drum, linkage, and handle to the user's hand and arm. The common model for these dynamics is a single mass connected to the user's desired position by a linear spring and damper; although second-order models perform well in other applications, the relevant dynamics of this complex system and the performance requirements for cancellation control must be closely examined before selecting an induced motion model.

5.6.1 Model Requirements

A cancellation controller will use the model $\hat{G}_{i,hf}$ for real-time simulation during telerobotic interactions. A particular model should thus be evaluated by its ability to capture motor movement under force feedback as a user attempts to hold the handle stationary, following the superposition assumption from Sec. 5.2. The model should be tested using force transients recorded from contact between the system's slave and environment. These signals contain frequencies from DC up to several hundred hertz, depending on the materials involved and the combined bandwidth of the tip-mounted force sensor and the feedback filter K_λ, if used. Model fidelity should be analyzed for a wide range of λ values, applying forces up to the maximum level that the master motor can generate. The importance of the system's response to these large-magnitude, high-frequency transients emphasizes the need to consider higher-order models, as internal master dynamics can create high-frequency resonances that strongly affect the induced motion response.

5.6.2 User Influence

While device dynamics are considered constant, the user's influence on the system varies over time. Each user's hand dynamics are unique, and the connection between hand and handle will naturally change with grip force and arm muscle co-contraction. To account for these biodynamics in the model, a time-varying compliant element can be placed between ground and the mass of the handle. Many researchers use spring-damper models to describe passive biodynamics [58, 61, 73], and my previous investigations have supported the efficacy

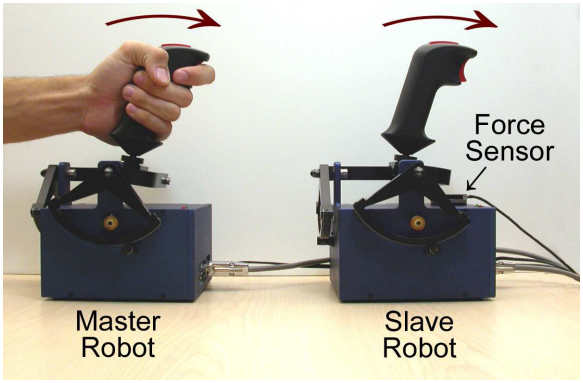

Figure 5.8: Single-axis position-force telerobotic testbed.

of such a user model in haptic interactions [80, 84]. The element's properties can be linked to an indication of the user's effective impedance such as a grip force measurement [15, 124] or an EMG reading [50, 118]. Most likely, the relationship between the sensed quantity and the hand's dynamic behavior will need to be calibrated for each user [84]. A general time-invariant model can be used only if variations in the user's dynamics are relatively small, such as if the human is required to maintain a constant grip on the handle.

In addition to his or her passive biodynamic influence on the master system, the user can actively respond to low-frequency feedback signals. At low frequency, the entire master moves as one object, pushing back against the user's hand when the slave sustains contact with an obstacle. The user's arm and the master's structure naturally deflect under such loads, but these deflections do not affect the user's ability to control the slave position under visual feedback. The user effectively removes low-frequency induced motion from the slave command via active compensation, applying additional hand forces and adjusting hand position at frequencies from steady state to about eight hertz. To avoid interfering with this normal, stable process, the model used in a cancellation controller must significantly underestimate low-frequency induced motion while preserving accurate behavior at high frequency. As seen in (5.14), overestimating the total response of the user-master system creates positive feedback, which can destabilize the control loop and should be avoided.

5.7 Experimental Evaluation

The modeling and control techniques presented in this chapter were developed through implementation on a one-dof telerobotic testbed. The system's master and slave are a

pair of Impulse Engine 2000's, high-quality force-feedback joysticks produced by Immersion Corporation. As shown in Figure 5.8, the forward-backward axis of each device includes a DC motor connected to the handle via a cable drive, drum, and mechanical linkage. The rectangular joystick bases contain the motors and their respective current amplifiers. An ATI Mini40 force sensor is located beneath the front transmission element of the slave joystick so that the contact force between it and the slave can be measured throughout an interaction. The hardware is controlled by a desktop computer running RTAI Linux, a hard-real-time open-source operating system, and all device control software was written in C. A five kilohertz servo loop reads the sensors, computes control forces, and outputs current commands to the amplifiers. This section presents experimental results from modeling and canceling this system's induced master motion.

5.7.1 Modeling Induced Master Motion

The dynamics of the master's induced motion pathway were identified via successive isolation, as described in Section 3.5.2. This process yielded the sixth-order nonlinear model described in Section 3.5.4, including hysteretic linkage stiffness and Coulomb friction at the motor and drum. During initial investigation, a linear stiffness model was found to adequately describe the user's behavior in the tested high-frequency range for moderate grip force values. Following the arguments of Section 5.6.2, the model's large deflections at DC were limited by adopting a combined linear and quintic user stiffness model. The user's quintic stiffness, q_u, was tuned to allow only two millimeters of deflection at maximum feedback force, giving $q_u = 3.125 * 10^{13}$ N/m^5. This updated user model thereby avoids overestimation of induced master motion at steady state, enabling the full model to capture the high-frequency, low-magnitude resonance that destabilizes teleoperation. Although this model did not incorporate variations in the user's hand impedance during an interaction, it could easily be extended to do so in future work.

With all of its parameters identified, tests were performed on the dynamic model's ability to predict the master's time response during haptic interactions. Several pre-recorded impact forces were displayed to a user via open-loop output. During these tests, the user held the handle of the master device, moving it forward as though tapping on a remote environment. The force profile was displayed when the user crossed a position threshold, and the resulting master motion, x_m, was recorded. The model was subjected to the same feedback force, F_f, and its prediction of high-frequency induced master motion, $\hat{x}_{mi,hf}$, was

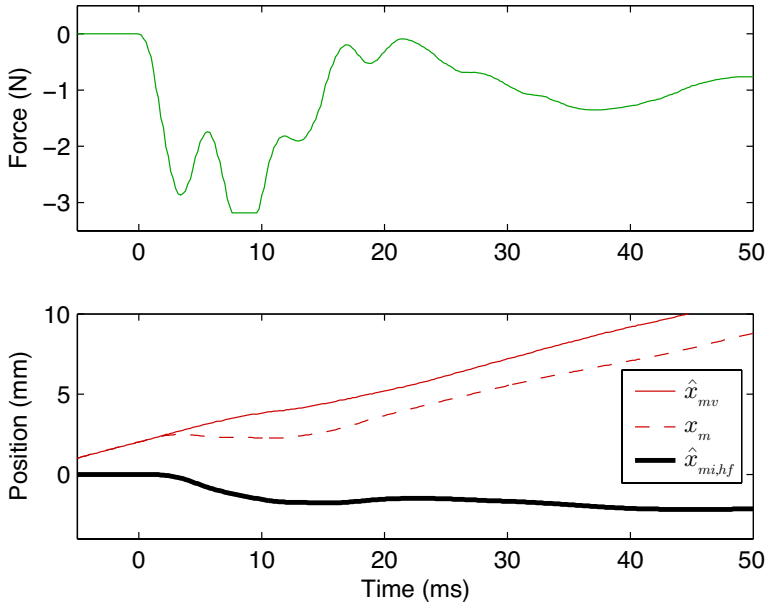

Figure 5.9: Full model prediction of induced master motion during open-loop display of a pre-recorded force profile.

recorded. The cancellation approach's estimate of intended user motion was formulated for each of the tests as

$$\hat{x}_{mv} = x_m - \hat{x}_{mi,hf}. \tag{5.20}$$

Open-loop testing indicates that the model adeptly captures the system's transient response to force feedback; results from a sample test are shown in Figure 5.9. The induced motion prediction matches the shape and magnitude of the haptic interface's observed deviations for a range of force profiles, feedback gains, and incoming velocities. The resulting estimates of intended motion are smooth, nearly devoid of evidence of the applied force transient, supporting our assumption of superposition. The model's ability to estimate induced master motion during open-loop force display bodes well for its use in a closed-loop controller.

5.7.2 Canceling Induced Master Motion

The one-degree-of-freedom telerobotic system in Figure 5.8 was used to test the cancellation approach with the identified model. Position-force control was implemented on the system, using proportional and derivative feedback on the slave. The full nonlinear master model

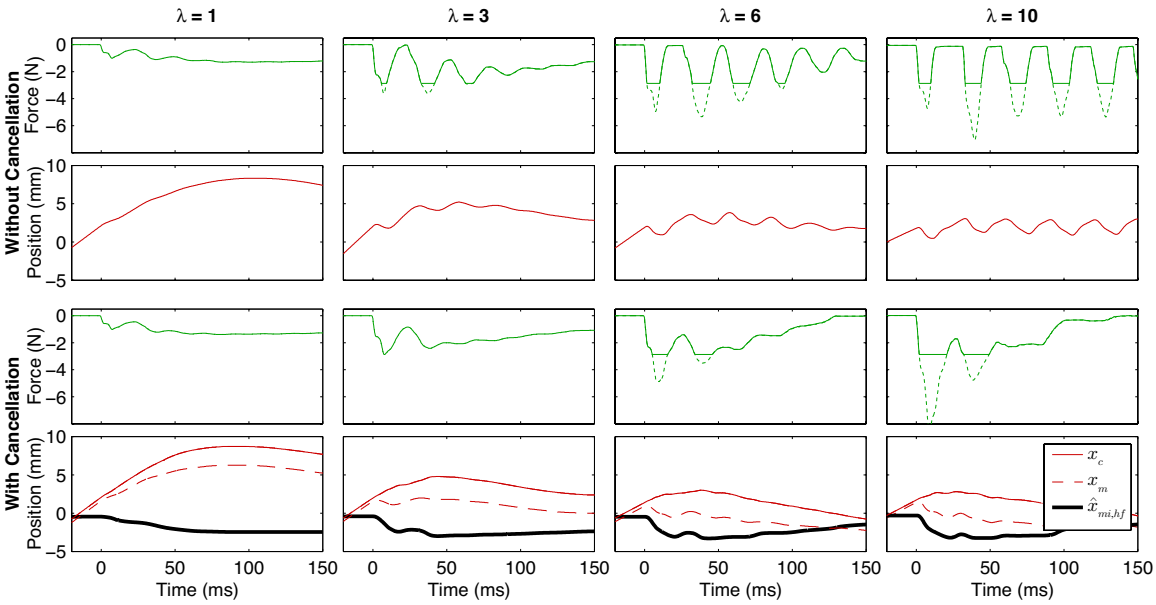

Figure 5.10: Force feedback, F_f, and slave position command, x_c, with and without cancellation for four λ values, keeping $\mu = 1$.

was simulated in real time using forward Euler integration with a time step of 40 μs. The performance of this simple teleoperator was tested with and without cancellation, setting the force-feedback gain, λ, to various values, and keeping the position scale, μ, at unity. In these tests, the user tapped repeatedly on the environment, as though trying to explore its surface.

Figure 5.10 shows a single representative contact for each of the eight testing conditions, illustrating the saturated force feedback signal, F_f, and commanded slave position, x_c. The bottom row of figures, which include cancellation, also show the measured master position, x_m, and the model's estimate of high-frequency induced master motion, $\hat{x}_{mi,hf}$. As seen in the first column of Figure 5.10, contact with the environment at $\lambda = 1$ produces nearly identical behavior with and without cancellation. The force transients decay in about 50 ms, inducing only a small amount of high-frequency master motion. Although stable, this configuration does not provide the user with significant force cues.

When the feedback gain in increased to $\lambda = 3$, as shown in the second column of Figure 5.10, the two controllers behave very differently. Without cancellation, the master experiences significant induced motion and the contact transient is prolonged to over 100 ms. The user experiences this effect as a slight buzzing on impact, which is completely removed

by cancellation. The compensated system accurately predicts the induced master motion and prevents it from contaminating the slave's position command. By $\lambda = 6$, contact without cancellation creates highly oscillatory force feedback, as the slave makes and breaks contact with the environment several times. Adding cancellation provides a smooth estimate of the user's intended path for the slave, enabling stable contact under amplified force feedback. Increasing the gain even further, to $\lambda = 10$, brings about contact instability in the uncompensated system. The violent shaking of the master mechanism is prevented by including cancellation in the controller, keeping the force's settling time at approximately fifty milliseconds. As hoped, cancellation enables the user to receive amplified environment force feedback that is not distorted by the dynamics of the telerobotic system's internal control loop. The higher λ values achievable with the CIMM approach give the user more choice in how to configure the system for different tasks; increasing the force reflection ratio creates stronger high-frequency accelerations at the master handle, which may give the user a more vivid impression of the remote environment's response to contact.

The strategy of canceling induced master motion was also examined through a series of blind user tests. An individual used an earlier version of the joystick teleoperation system to tap repeatedly on the hard surface of the force sensor. The forward position gain, μ, was set to one, and the force feedback gain, λ, was varied inversely with incoming velocity to keep the total magnitude of the haptic feedback within the device's limits. The system randomly assigned whether each trial included cancellation, and the user was not aware of these controller modifications. The force contact transient was recorded for each tap, and its settling time was determined via post-processing that was again blind to the controller type.

Figure 5.11 plots the measured settling times for all trials, including rough fit lines for those without cancellation and those with cancellation. For low force reflection ratios, both controllers are stable, but cancellation keeps the settling time of the transient to a lower value, closer to that of the slave's natural impact. Without cancellation, the system cannot sustain a force reflection ratio above about 2.5, while cancellation enables stable contact up to a gain of 4.5. The lower settling times afforded by cancellation indicate a more stable loop and less feedback distortion for the user. The λ values at which instability occurred differ somewhat from those shown in Figure 5.10 because the transmissions of both master and slave contained significant backlash that had not yet been corrected. Additionally, the tests shown in Figure 5.11 were performed with a simpler, less-accurate model of the

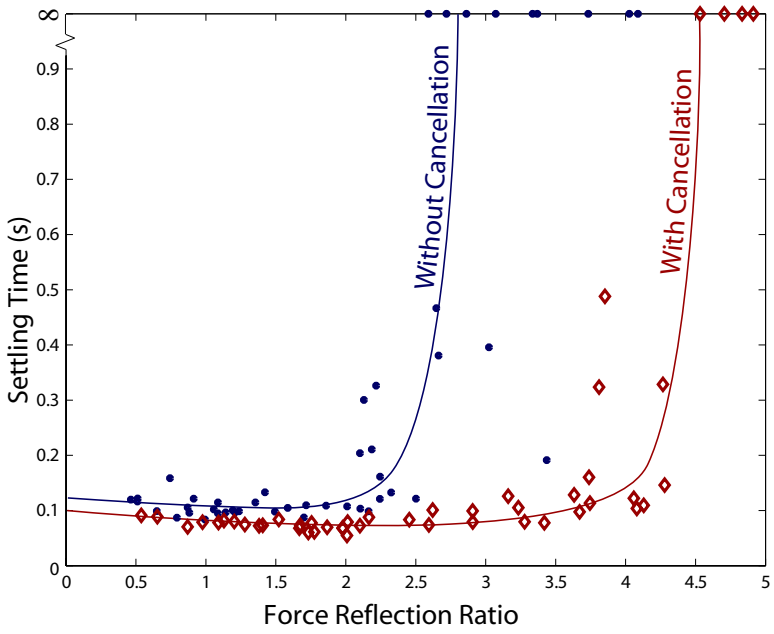

Figure 5.11: Settling time of the slave's contact force transient without and with cancellation of high-frequency induced master motion.

human-master dynamics. Despite these limitations, the CIMM approach was sufficient to enable over a two-fold increase in the stable force reflection gain, and the feel of the interface was significantly improved.

5.8 Summary

Force-reflecting teleoperation has historically been plagued by contact instability, preventing the use of high feedback gains and leaving the user with faint haptic cues. This phenomenon can be traced to the dynamics of the master device, which must simultaneously measure the user's position command and apply haptic feedback. The position-force architecture inadvertently creates an internal controller loop, as force feedback induces motion of the master mechanism that is not intended as a position command. Large forward and feedback gains allow high-frequency oscillations to resonate in this internal controller loop, prolonging contact transients and ultimately causing system instability.

The controller's destabilizing inner loop can be severed by removing induced master motion from the slave's position command. The relationship between feedback force and

master motion can be characterized by careful examination of device dynamics. The resulting model is simulated in real time during telerobotic interactions, and its output is subtracted from the measured master position to provide an estimate of user intention. This strategy was demonstrated on a one-dof master, using a nonlinear sixth-order model that successfully captures the transient response of the system as held by a user.

Model-based cancellation was found to stabilize contact on the one-dof telerobotic testbed, preventing the violent oscillations that typically occur at high force reflection gains. The CIMM approach attenuates the control loop's induced motion pathway and allows the system to provide stronger, more authentic force feedback to the user. With a carefully derived model, the joystick system enabled stable tapping up to a feedback gain of 15 with approximately constant transient decay time. Further work is necessary to determine the relationship between model fidelity and system performance, as well as the robustness of this approach to model uncertainty.

The perspective of induced master motion also informs the process of master mechanism design. All impedance-type devices will allow some level of induced motion, so care should be taken to consider these dynamics. Generally, minimizing compliance in the device transmission will limit the induced motor motion pathway and provide better performance; suggestions from [146] may prove useful in this process. Because it relies on accurate modeling of the haptic interface's dynamics, canceling induced master motion may provide more significant performance benefits on systems that contain a relatively soft connecting element, like the joystick's linkage. Having a soft connection between handle and motor reduces the influence of the user's time-varying properties on the dynamics that govern induced master motion, making them easier to estimate with an LTI model. I hope that the methods of modeling and canceling induced master motion will add to the field's existing set of control techniques and enable future teleoperators to portray haptic interactions more authentically, allowing users to telerobotically feel the full spectrum of haptic feedback that is available during real interactions.

Chapter 6

Conclusion

Haptic interfaces perform the marvelous feat of allowing humans to touch objects that are far away or that exist only as computerized models. Present telerobotic systems enable surgeons to repair the human heart through tiny incisions [129] and let astronauts maneuver hefty payloads from inside the International Space Station [95]. Current haptic interfaces also let artists sculpt virtual clay [97] and give veterinarians-in-training the opportunity to practice palpating a digital simulation of the bovine reproductive tract [7]. Despite the great advances these technologies have seen, one major factor that still limits their usefulness is the remarkably poor feel of contact with hard or textured objects.

Soft, indistinct haptic feedback can be traced to the control methodologies commonly employed on such systems, as they fail to create appropriate high-frequency accelerations at the user's fingertips. A closed-loop position-based algorithm can make an autonomous robot arm adeptly track a programmed trajectory, but it cannot enable an impedance-type haptic device to portray impacts with stiff surfaces, regardless of whether the position signal stems from a remote robot's movement or a virtual environment's geometry. Event-based feedback can improve the crispness of virtual surfaces, but previous efforts have required extensive tuning of transient parameters by feel rather than by deterministic guidelines. For teleoperation, the position-force architecture captures the vibratory details of remote contact, but its feedback force must almost always be attenuated to avoid contact instability.

This thesis confronts the limitations of present haptic feedback algorithms by focusing on the high-frequency internal dynamics of the haptic interface, the system that connects the user to the controller. Most researchers assume that these devices provide an ideal rigid connection between the user's hand and the system's sensors and actuators, but real systems

exhibit significant structural compliance, bearing friction, other nonlinear behaviors, and variations over time. Only by understanding the influence of its electrical, mechanical, and biomechanical elements can a haptic interface designer hope to accurately portray the nuanced experience of touching real objects. The earlier chapters of this dissertation aimed to provide a technical basis for characterizing and controlling the high-frequency dynamics of haptic interfaces to improve the feel of hard and textured objects; the chief contributions of this work are listed below in Section 6.1, and Section 6.2 provides relevant guidelines for improving the performance of existing haptic interfaces. The discussion concludes in Section 6.3 with suggestions for future research on this topic.

6.1 Contributions

The primary contributions of this thesis are the following:

- An efficient approach for characterizing the high-frequency (10 to 1000 hertz) dynamics of haptic interfaces using empirical transfer function estimates: comprehensive evaluation facilitates a frequency-domain understanding of the relationship between motor current command and handle acceleration or motor movement while a haptic interface is held by a user, and it can be extended to capture the influence of time-varying parameters like user grip force.

- A thorough procedure for obtaining a physically-based lumped-parameter model of a haptic interface: successive isolation combines a variety of time- and frequency-domain identification techniques to estimate inertial, compliant, and dissipative parameters, working through each element of the dynamic chain in turn, from the computer to the user's hand.

- A method for modeling rate-independent hysteresis in compliant dynamic elements: close alignment between simulation and observation can be achieved by switching the element's force output between two offset stiffness lines based on the sign of its deflection change and smoothing this signal with a tunable spatial filter.

- A teleoperation feedback architecture that allows the user to feel the precise details of the slave robot's contact with rough and stiff objects in the remote environment: high-frequency acceleration matching uses an accelerometer on the slave's end effector

to record contact vibrations and recreates these accelerations at the user's hand in real time by inverting a model of the haptic interface's dynamics and overlaying this high-frequency feedback channel with traditional position-position control.

- A deterministic method for making virtual objects of various materials feel realistically hard: high-frequency acceleration matching enables the event-based portrayal of appropriate fingertip accelerations during contact with virtual objects by conditioning a set of pre-recorded acceleration transients with the inverse of a model of the haptic interface's dynamics.

- A control technique that stabilizes strong high-frequency haptic feedback like the signals available in position-force teleoperation: canceling induced master motion uses a model of the haptic interface's dynamics to estimate the motor motion that will result from applied feedback commands and remove it from the system's estimation of the user's intended hand motion.

The developed modeling techniques can be used on their own to better understand the behavior of existing haptic interfaces. High-frequency acceleration matching and/or cancellation of induced master motion can then be added to a system's controller to improve the feel of contact with hard and textured objects. These high-frequency feedback methods methods can be combined with the many previously developed strategies for improving the performance and stability of haptic interfaces, thereby achieving new levels of realism and broadening the applicability of such systems.

6.2 Implications

This research delved deep into the inner workings of several impedance-type haptic interfaces, each of which was a unique combination of computational, electrical, and mechanical elements. Over the course of these investigations, I have formulated a set of guidelines for achieving top performance from such systems, especially in the area of high-frequency haptic feedback. These practical implications may aid those who seek to build on this work in the future.

The computer coordinates and controls the behavior of the haptic interface. In order to run the servo loop at rates far above the human perception bandwidth of one kilohertz,

the computer system needs to be able to perform tasks such as encoder reads and digital-to-analog writes in a very short amount of time. For this reason, control cards that operate on the computer's PCI bus are far preferred to those that utilize the slower ISA bus or those that use serial communications protocols such as USB. The accuracy with which the computer executes these tasks must be very high as well, so that the program knows the amount of time elapsed since the prior step; this knowledge is crucial for computing signal derivatives such as velocity and also for running a discretized dynamic model in real time. A hard real-time operating system such as RTAI Linux adeptly provides these important computational features.

The electrical connections between the computer and the haptic mechanism should also be carefully configured to ensure high performance. Though utilized in most commercial haptic interfaces, PWM current amplifiers are not necessarily the best choice for generating high-frequency haptic feedback. They are often conservatively tuned to have low bandwidth, and their cut-off frequency depends on the inductance of the motor. Their fast switching excessively heats the motor and injects high-frequency electrical noise into neighboring sensor signals. Finally, they commonly exhibit non linearities such as cross-over distortion, which are hard to model and invert. High-performance PWM amplifiers can overcome these disadvantages, but linear current amplifiers are often more appropriate for use with haptic interfaces, as they can provide clean, high-bandwidth output. Analog input lines like those from accelerometers should be shielded nonetheless, and the corresponding sensor power lines should be heavily low-pass filtered to remove ambient electrical noise. Finally, power for optical encoders may need to be secured from an external source rather than the computer's motherboard in order to ensure accurate reads during the very fast motor motion that can occur at resonant peaks during frequency-domain identification.

A haptic interface's mechanical components are the last class of elements that strongly influence its high-frequency behavior. All physical joints and connections should be checked for backlash, and this play should be removed by remachining, shimming, or the application of additional fasteners. Cable tension should be checked and adjusted regularly, and worn cables should be replaced. Bearing friction should be evaluated, and faulty components should be replaced. Soft mechanical elements in the chain should be stiffened as much as possible to facilitate the transmission of high-frequency signals from the motor to the user's hand and to minimize the interface's induced master motion. If a lumped-parameter model of the interface has been derived through successive isolation, the effects of component

changes can be simulated, and the design can be fine-tuned accordingly. Generally, it is desirable for a magnitude response not to contain very high peaks or very deep valleys, as such behaviors can change significantly with small parameter variations, which makes it difficult to compensate for them.

Incorporating all of these suggestions will yield a haptic interface that is well equipped for the high-frequency modeling and control methods presented in this research. It will not behave ideally, as most researchers assume, but it will perform well. A good model of its dynamics should be obtained through the steps outlined in Chapter 3, including characterization of the influence of user grip force. Such a model allows the designer to consider the full span of possible dynamic behaviors for the haptic interface, ranging from its independent mass-like response to the more complex behavior it exhibits when held by a user. Once a good dynamic model has been generated, the feel of hard and textured contact can be improved through high-frequency acceleration matching, as described in Chapter 4. If contact with hard remote or virtual objects is unstable or vibratory, Chapter 5's method of cancellation of induced master motion can be added to the system.

6.3 Future Work

This dissertation establishes the feasibility and value of high-frequency acceleration matching and cancellation of high-frequency induced master motion, both of which are enabled by careful characterization of the internal dynamics of haptic interfaces. As demonstrated on several one-degree-of-freedom systems, these measures improve the feel of hard or textured contacts over position-based feedback, and they provide greater stability than that afforded by high-magnitude position-force control. With initial development and testing complete, more research will be needed to facilitate widespread application of these ideas, and further investigations will be inspired by the current and future findings.

The first natural extension of the present research is to apply it to a three-dimensional system. A good model would need to be obtained for each axis of a haptic interface; initial investigations revealed that the base and elbow of the studied Phantom exhibit similar, though not identical, dynamics to those of the fully characterized shoulder. The dynamic relationships between the axes would also need to be investigated to determine whether a fully decoupled assumption is appropriate. It is likely that motor position response of the base would be found to be independent from that of the shoulder and the elbow, but the

shoulder and elbow's position responses might exhibit some coupling, as they are driven by the same cable. Lower-frequency handle accelerations would probably be determined to be independent, but the three axes of high-frequency handle accelerations would probably be somewhat coupled. Determining an appropriate method for inverting these coupled dynamics is an interesting, important challenge.

Another exciting effort in this realm would be to obtain dynamic models for all widely available haptic interfaces. The models could be shared and compared between researchers via a central database to determine the suitability of different devices for various applications. Ideally, each manufacturer would create the dynamic models for their systems and share these with the user community, in the same way that geometric models are presently disseminated. In support of this effort, the characterization process would surely need to be streamlined. Development of an automatic fitting method of LTI models to ETFE data would be very useful, and it could help determine an appropriate metric for quantifying the fit of frequency-domain models. The best models for this purpose would the physically-based lumped-parameter ones that are obtained through successive isolation; manufacturers are in a unique position to help create these models, as they know the specifications and geometry of every element in the device. Hopefully, an attentiveness to high-frequency performance will inspire the next generation of haptic interfaces to be designed with such objectives in mind. Over time, we will also discover the level of detail required from dynamic models of haptic interfaces; the methods presented here yield very precise predictions at the expense of complexity, and simpler models may eventually be found to work well in certain situations.

Another exciting area for future work is to help develop a greater understanding of the human sense of touch. High-frequency acceleration matching allows the haptic interface designer to create a wide range of fingertip accelerations, including ones that are based on physical interactions and ones that are generated artificially. With the power to vary magnitude, frequency, shape, and duration, a psychophysicist can begin to determine which aspects of a contact transient are most salient to users. These findings can be used to guide the design of high-frequency feedback in haptic interfaces and may also shed light on the functionality of the human tactile sensory system. For example, such investigations could be used to explore the user's perception of the rendered slave tip acceleration during teleoperation with HFAM; perhaps filtering or frequency-shifting this signal would improve the realism of the interaction over the straight matching that I have presently proposed.

Another interesting effort in this area would be to map out the range of virtual hard contacts that can be portrayed with a typical haptic interface and study the connection between distortion of the acceleration transient and the haptic impression that the user forms of the contact.

Finally, this work constitutes the foundation of a long-ranging research approach that I call haptography: improving the authenticity of remote and virtual touch by carefully matching haptic feedback to the real experiences being emulated. Despite its ubiquitous importance in human life, we have relatively little practical knowledge about the sense of touch, and we lack a formal method for describing haptic experiences. In contrast, consider the human mastery of visual stimuli; to create a lifelike image of an observed scene, we no longer need to start with a blank canvas, painting each element by hand. Instead, we use a sophisticated measuring system (a camera lens) to control the desired light pattern and project it onto a sensitive medium (either film or a digital image sensor). The latent image is then converted to a storable, portable record (a negative or an image file) using chemical or signal processing. Finally, we create realistic copies of the original stimulus by printing the picture on paper or displaying it on an imaging screen. My future research will seek to understand and control haptic interactions to this same level of excellence, building on this base of knowledge in high-frequency haptic feedback for hard and textured contacts.

Bibliography

[1] Richard J. Adams, Manuel R. Moreyra, and Blake Hannaford. Stability and performance of haptic displays: Theory and experiments. In *Proc. ASME International Mechanical Engineering Congress and Exhibition*, pages 227–234, November 1998.

[2] Marco Agus, Andrea Giachetti, Enrico Gobbetti, Gianluigi Zanetti, and Antonio Zorcolo. Real-time haptic and visual simulation of bone dissection. *Presence: Teleoperators and Virtual Environments*, 12(1):110–122, February 2003.

[3] Robert J. Anderson and Mark W. Spong. Bilateral control of teleoperators with time delay. *IEEE Transactions on Automatic Control*, 34(5):494–501, May 1989.

[4] Paolo Arcara and Claudio Melchiorri. A comparison of control schemes for teleoperation with time delay. In *Proc. IFAC Conference on Telematics Applications in Automation and Robotics*, pages 1–6, July 2001.

[5] Roland Arsenault and Colin Ware. Eye-hand co-ordination with force feedback. In *Proc. SIGCHI Conference on Human Factors in Computing Systems*, pages 408–414, April 2000.

[6] Jonathan F. Ashmore. Mechanoreception. In Antonio Borsellino, Luigi Cervetto, and Vincent Torre, editors, *Sensory Transduction*, volume 194 of *NATO ASI Series A: Life Sciences*, pages 25–50. Plenum Press, New York, 1990.

[7] Sarah Baillie, Dominic J. Mellor, Stephen A. Brewster, and Steward W. J. Reid. Integrating a bovine rectal palpation simulator into an undergraduate veterinary curriculum. *Journal of Veterinary Medical Education*, 32(1):79–85, Spring 2005.

[8] Cagatay Basdogan, Chih-Hao Ho, and Mandayam A. Srinivasan. Virtual environments in medical training: Graphical and haptic simulation of laparoscopic common

bile duct exploration. *IEEE/ASME Transactions on Mechatronics*, 6(3):269–285, September 2001.

[9] Cagatay Basdogan and Mandayam A. Srinivasan. Haptic rendering in virtual environments. In Kay M. Stanney, editor, *Handbook of Virtual Environments: Design, Implementation, and Applications*, chapter 6, pages 117–134. Lawrence Erlbaum Associates, London, 2002.

[10] Jonathan Bell, Stanley Bolanowski, and Mark H. Holmes. The structure and function of Pacinian corpuscles: A review. *Progress in Neurobiology*, 42(1):79–128, January 1994.

[11] Laura Bianchi and Monica Driscoll. The molecular basis of touch sensation as modeled in Caenorhabditis elegans. In Stephan Frings and Jonathan Bradley, editors, *Transduction Channels in Sensory Cells*, chapter 1, pages 1–29. Wiley, New York, October 2004.

[12] A. J. Brisben, S. S. Hsiao, and K. O. Johnson. Detection of vibration transmitted through an object grasped in the hand. *Journal of Neurophysiology*, 81(4):1548–1558, April 1999.

[13] Yonghong Bu, R. W. Daniel, and P. R. McAree. Stability analysis of force reflecting telerobotic systems. In *Proc. IEEE International Conference on Intelligent Robots and Systems*, volume 3, pages 1374–1379, November 1996.

[14] Grigore C. Burdea. *Force and Touch Feedback for Virtual Reality*. John Wiley and Sons, New York, 1996.

[15] Lage Burström. The influence of biodynamic factors on the mechanical impedance of the hand and arm. *International Archives of Occupational and Environmental Health*, 69(6):437–446, June 1997.

[16] Pietro Buttolo, Darwei Kung, and Blake Hannaford. Manipulation in real, virtual and remote environments. In *Proc. IEEE International Conference on Systems, Man, and Cybernetics*, pages 4656–4661, October 1995.

[17] Guy-Bernard Cadière, Jacques Himpens, Olivier Germay, Rachel Izizaw, Michel Degueldre, Jean Vandromme, Elie Capelluto, and Jean Bruyns. Feasibility of robotic

laparoscopic surgery: 146 cases. *World Journal of Surgery*, 25(11):1467–1477, November 2001.

[18] Murat Cenk Çavuşoğlu, David Feygin, and Frank Tendick. A critical study of the mechanical and electrical properties of the PHANToM haptic interface and improvements for high-end performance control. *Presence: Teleoperators and Virtual Environments*, 11(6):555–568, December 2002.

[19] Murat Cenk Çavuşoğlu, Alana Sherman, and Frank Tendick. Bilateral controller design for telemanipulation in soft environments. *IEEE Transactions on Robotics and Automation*, 18(4):641–647, August 2002.

[20] L. Chang, R. M. Satava, C. A. Pellegrini, and M. N. Sinanan. Robotic surgery: Identifying the learning curve through objective measurement of skill. *Surgical Endoscopy*, 17(11):1744–1748, November 2003.

[21] Robert R. Christensen, John M. Hollerbach, Yangming Xu, and Sanford G. Meek. Inertial-force feedback for the treadport locomotion interface. *Presence: Teleoperators and Virtual Environments*, 9(1):1–14, February 2000.

[22] J. Edward Colgate. Robust impedance shaping telemanipulation. *IEEE Transactions on Robotics and Automation*, 9(4):374–384, August 1993.

[23] J. Edward Colgate, Paul E. Grafing, Michael C. Stanley, and Gerd Schenkel. Implementation of stiff virtual walls in force-reflecting interfaces. In *Proc. IEEE Virtual Reality Annual International Symposium*, pages 202–208, September 1993.

[24] Daniela Constantinescu, Septimiu E. Salcudean, and Elizabeth A. Croft. Haptic rendering of rigid body contacts using impulsive and penalty forces. *IEEE Transactions on Robotics*, 21(3):309–323, June 2005.

[25] Renato Costi, Jacques Himpens, Jean Bruyns, and Guy Bernard Cadière. Robotic fundoplication: From theoretic advantages to real problems. *Journal of the American College of Surgeons*, 197(3):500–507, September 2003.

[26] Stéphane Cotin, Hervé Delingette, and Nicholas Ayache. A hybrid elastic model for real-time cutting, deformations, and force feedback for surgery training and simulation. *The Visual Computer*, 16(8):437–452, December 2000.

[27] James Cremer, Joseph Kearney, and Yiannis Papelis. Driving simulation: Challenges for VR technology. *IEEE Computer Graphics and Applications*, 16(5):16–20, September 1996.

[28] R. W. Daniel and P. R. McAree. Fundamental limits of performance for force reflecting teleoperation. *International Journal of Robotics Research*, 17(8):811–830, August 1998.

[29] Thomas Debus, Theresia Becker, Pierre Dupont, Tae-Jeong Jang, and Robert D. Howe. Multichannel vibrotactile display for sensory substitution during teleoperation. In Matthew R. Stein, editor, *Proc. SPIE Telemanipulator and Telepresence Technologies VIII*, volume 4570, pages 42–49, February 2002.

[30] Jack Tigh Dennerlein, Paul A. Millman, and Robert D. Howe. Vibrotactile feedback for industrial telemanipulators. In *Proc. ASME Dynamic Systems and Control Division*, volume 61, pages 189–195, November 1997.

[31] R. Q. Van der Linde, P. Lammertse, E. Frederiksen, and B. Ruiter. The HapticMaster, a new high-performance haptic interface. In S. A. Wall, B. Riedel., A. Crossan, and M. R. McGee, editors, *Proc. EuroHaptics Conference*, pages 1–5, July 2002.

[32] Nicola Diolaiti, Günter Niemeyer, Federico Barbagli, and J. Kenneth Salisbury. Stability of haptic rendering: Discretization, quantization, time delay, and Coulomb effects. *IEEE Transactions on Robotics*, 22(2):256–268, April 2006.

[33] Nicola Diolaiti, Günter Niemeyer, Federico Barbagli, J. Kenneth Salisbury, and Claudio Melchiorri. The effect of quantization and Coulomb friction on the stability of haptic rendering. In *Proc. IEEE World Haptics Conference*, March 2005.

[34] Encarta. World English Dictionary. Microsoft Corporation, 2006.

[35] Volkmar Falk, Anno Diegeler, Thomas Walther, Jurgen Banusch, Jan Brucerius, Jorg Raumans, Rudiger Autschbach, and Friedrich W. Mohr. Total endoscopic computer enhanced coronary artery bypass grafting. *European Journal of Cardio-thoracic Surgery*, 17:38–45, January 2000.

[36] Eric L. Faulring. *The Cobotic Hand Controller: Design, Control and Analysis of a Novel Haptic Display.* Ph.D. Thesis, Northwestern University, Evanston, IL, 2005.

[37] Eric L. Faulring, Kevin M. Lynch, J. Edward Colgate, and Michael A. Peshkin. Haptic interaction with constrained dynamic systems. In *Proc. IEEE International Conference on Robotics and Automation*, pages 2458–2464, April 2005.

[38] Jonathan Fiene, Katherine J. Kuchenbecker, and Günter Niemeyer. Event-based haptic tapping with grip force compensation. In *Proc. IEEE Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 117–123, March 2006.

[39] Kevin B. Fite, John E. Speich, and Michael Goldfarb. Transparency and stability robustness in two-channel bilateral telemanipulation. *ASME Journal of Dynamic Systems, Measurement, and Control*, 123(3):400–407, September 2001.

[40] FlexiForce, a division of Tekskan, http://www.tekscan.com/flexiforce.html.

[41] Jason P. Fritz and Kenneth E. Barner. Stochastic models for haptic texture. In Matthew R. Stein, editor, *Proc. SPIE Telemanipulator and Telepresence Technologies III*, volume 2901, pages 34–44, November 1996.

[42] Frank A. Geldard. *The Human Senses*. Wiley, New York, 2nd edition, 1972.

[43] R. Brent Gillespie and Mark Cutkosky. Stable user-specific rendering of the virtual wall. In *Proc. ASME Dynamic Systems and Control Division*, volume 58, pages 397–406, November 1996.

[44] R. Brent Gillespie, Christopher Hasser, and Philip Tang. Cancellation of feedthrough dynamics using a force-reflecting joystick. In *Proc. ASME Dynamic Systems and Control Division*, volume 67, pages 319–326, November 1999.

[45] Raymond C. Goertz. Fundamentals of general-purpose remote manipulators. *Nucleonics*, 10(11):36–42, November 1952.

[46] Raymond C. Goertz. Mechanical master-slave manipulator. *Nucleonics*, 12(11):45–46, November 1954.

[47] Raymond C. Goertz and Frank Bevilacqua. A force-reflecting positional servomechanism. *Nucleonics*, 10(11):43–45, November 1952.

[48] Raymond C. Goertz and W. M. Thompson. Electronically controlled manipulator. *Nucleonics*, 12(11):46–47, November 1954.

[49] Michael Goldfarb. Dimensional analysis and selective distortion in scaled bilateral telemanipulation. In *Proc. IEEE International Conference on Robotics and Automation*, volume 2, pages 1609–1614, May 1998.

[50] Paul L. Gribble, Lucy I. Mullin, Nicholas Cothros, and Andrew Mattar. Role of cocontraction in arm movement accuracy. *Journal of Neurophysiology*, 89(5):2396–2405, May 2003.

[51] Weston B. Griffin, William R. Provancher, and Mark R. Cutkosky. Feedback strategies for telemanipulation with shared control of object handling forces. *Presence: Teleoperators and Virtual Environments*, 14(6):720–731, December 2005.

[52] Gary S. Guthart and J. Kenneth Salisbury, Jr. The Intuitive^{TM} telesurgery system: Overview and application. In *Proc. IEEE International Conference on Robotics and Automation*, volume 1, pages 618–621, April 2000.

[53] Aram Z. Hajian and Robert D. Howe. Identification of the mechanical impedance at the human finger tip. *Journal of Biomechanical Engineering*, 119(1):109–114, February 1997.

[54] Blake Hannaford. A design framework for teleoperators with kinesthetic feedback. *IEEE Transactions on Robotics and Automation*, 5(4):426–434, August 1989.

[55] Blake Hannaford and Robert Anderson. Experimental and simulation studies of hard contact in force reflecting teleoperation. In *Proc. IEEE International Conference on Robotics and Automation*, volume 1, pages 584–589, April 1988.

[56] Blake Hannaford, Laurie Wood, Douglas A. McAffee, and Haya Zak. Performance evaluation of a six-axis generalized force-reflecting teleoperator. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(3):620–633, May/June 1991.

[57] Keyvan Hashtrudi-Zaad and Septimiu E. Salcudean. Analysis of control architectures for teleoperation systems with impedance/admittance master and slave manipulators. *International Journal of Robotics Research*, 20(6):419–445, June 2001.

[58] Christopher J. Hasser and Mark R. Cutkosky. System identification of the human grasping a haptic knob. In *Proc. IEEE Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 171–180, March 2002.

[59] Graham S. Hawkes. Apparatus providing tactile feedback to operators of remotely controlled manipulators. United States Patent No. 4,655,673, April 7, 1987.

[60] Vincent Hayward and Brian Armstrong. A new computational model of friction applied to haptic rendering. In Peter I. Corke and James Trevelyan, editors, *Experimental Robotics VI*, volume 250 of *Lecture Notes in Control and Information Science*, pages 404–412. Springer, 2000.

[61] Neville Hogan. Controlling impedance at the man/machine interface. In *Proc. IEEE International Conference on Robotics and Automation*, volume 3, pages 1621–1631, May 1989.

[62] Robert D. Howe. A force-reflecting teleoperated hand system for the study of tactile sensing in precision manipulation. In *Proc. IEEE International Conference on Robotics and Automation*, volume 2, pages 1321–1326, May 1992.

[63] Robert D. Howe and Mark R. Cutkosky. Sensing skin acceleration for slip and texture perception. In *Proc. IEEE International Conference on Robotics and Automation*, volume 1, pages 145–150, May 1989.

[64] Jesse D. Hwang, Michael D. Williams, and Günter Niemeyer. Toward event-based haptics: Rendering contact using open-loop force pulses. In *Proc. IEEE Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 24–31, March 2004.

[65] J. M. Hyde and Mark R. Cutkosky. Controlling contact transition. *IEEE Control System's Magazine*, 14(1):25–30, February 1994.

[66] Immersion Corporation, http://www.immersion.com.

[67] Gunnar Jansson and Katarina Billberger. The PHANToM used without visual guidance. In *Proc. First PHANToM Users Research Symposium*, May 1999. Available at http://mbi.dkfz-heidelberg.de/purs99/.

[68] Roland S. Johansson. Receptive field sensitivity profile of mechanosensitive units innervating the glabrous skin of the human hand. *Brain Research*, 104(2):330–334, March 1976.

[69] Roland S. Johansson and Benoni B. Edin. Predictive feed-forward sensory control during grasping and manipulation in man. *Biomedical Research*, 14(Supp. 4):95–106, 1993.

[70] Roland S. Johansson, U. Landström, and R. Lundström. Response of mechanoreceptive afferent units in the glabrous skin of the human hand to sinusoidal skin displacements. *Brain Research*, 244(1):17–25, July 1982.

[71] Roland S. Johansson and Åke B. Vallbo. Tactile sensory coding in the glabrous skin of the human hand. *Trends in Neurosciences*, 6(2):27–32, 1983.

[72] Homayoon Kazerooni, Tsing-Iuan Tsay, and Karin Hollerbach. A controller design framework for telerobotic systems. *IEEE Transactions on Control Systems Technology*, 1(1):50–62, March 1993.

[73] R. E. Kearney and I. W. Hunter. System identification of human joint dynamics. *Critical Reviews in Biomedical Engineering*, 18(1):55–87, 1990.

[74] Masaya Kitagawa, Allison M. Okamura, Brian T. Bethea, Vincent L. Gott, and William A. Baumgartner. Analysis of suture manipulation forces for teleoperation with force feedback. In *Proc. Fifth International Conference on Medical Image Computing and Computer Assisted Intervention*, September 2002.

[75] Roberta Klatzky and Susan Lederman. The haptic identification of everyday life objects. In Yvette Hatwell, editor, *Touching for Knowing: Cognitive Psychology of Haptic Manual Perception*, chapter 7, pages 105–121. John Benjamins Publishing Company, Philadelphia, 2003.

[76] Dimitrios A. Kontarinis and Robert D. Howe. Tactile display of vibratory information in teleoperation and virtual environments. *Presence: Teleoperators and Virtual Environments*, 4(4):387–402, August 1995.

[77] Ken'ichi Koyanagi, Tomoko Morita, and Junji Furusho. Basic algorithm of controlling passive force display system with redundant brakes. In *Proc. IEEE International Conference on Robotics and Automation*, pages 1767–1772, April 2005.

[78] Katherine J. Kuchenbecker, Jonathan Fiene, and Günter Niemeyer. Event-based haptics and acceleration matching: Portraying and assessing the realism of contact. In *Proc. IEEE World Haptics Conference*, pages 381–387, March 2005.

[79] Katherine J. Kuchenbecker, Jonathan Fiene, and Günter Niemeyer. Improving contact realism through event-based haptic feedback. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):219–230, March/April 2006.

[80] Katherine J. Kuchenbecker and Günter Niemeyer. Canceling induced master motion in force-reflecting teleoperation. In *Proc. ASME International Mechanical Engineering Congress and Exposition*, volume 2, paper 60049, November 2004.

[81] Katherine J. Kuchenbecker and Günter Niemeyer. Modeling induced master motion in force-reflecting teleoperation. In *Proc. IEEE International Conference on Robotics and Automation*, pages 350–355, April 2005.

[82] Katherine J. Kuchenbecker and Günter Niemeyer. Improving telerobotic touch via high-frequency acceleration matching. In *Proc. IEEE International Conference on Robotics and Automation*, pages 3893–3898, May 2006.

[83] Katherine J. Kuchenbecker and Günter Niemeyer. Induced master motion in force-reflecting teleoperation. Accepted to ASME Journal of Dynamic Systems, Measurement and Control, 2006.

[84] Katherine J. Kuchenbecker, June Gyu Park, and Günter Niemeyer. Characterizing the human wrist for improved haptic interaction. In *Proc. ASME International Mechanical Engineering Congress and Exposition*, volume 2, paper 42017, November 2003.

[85] U. Kühnapfel, H. K. Çakmak, and H. Maaß. Endoscopic surgery training using virtual reality and deformable tissue simulation. *Computers and Graphics*, 24(5):671–682, October 2000.

[86] Robert H. LaMotte. Softness discrimination with a tool. *Journal of Neurophysiology*, 83(4):1777–1786, April 2000.

[87] C. Andrew Lawn and Blake Hannaford. Performance testing of passive communication and control in teleoperation with time delay. In *Proc. IEEE International Conference on Robotics and Automation*, volume 3, pages 776–781, May 1993.

[88] Dale A. Lawrence. Stability and transparency in bilateral teleoperation. *IEEE Transactions on Robotics and Automation*, 9(5):624–637, October 1993.

[89] Dale A. Lawrence, Lucy Y. Pao, Anne M. Dougherty, Mark A. Salada, and Yiannis Pavlou. Rate-hardness: A new performance metric for haptic interfaces. *IEEE Transactions on Robotics and Automation*, 16(4):357–371, August 2000.

[90] Dale A. Lawrence, Lucy Y. Pao, Mark A. Salada, and Anne M. Dougherty. Quantitative experimental analysis of transparency and stability in haptic interfaces. In *Proc. ASME Dynamic Systems and Control Division*, volume 58, November 1996.

[91] Hyoung-Ki Lee and Myung Jin Chung. Adaptive controller of a master-slave system for transparent teleoperation. *Journal of Robotic Systems*, 15(8):465–475, August 1998.

[92] Jack M. Loomis and Susan J. Lederman. Tactual perception. In Kenneth R. Boff, Lloyd Kaufman, and James P. Thomas, editors, *Handbook of Perception and Human Performance*, volume II: Cognitive Processes and Performance, chapter 31. John Wiley and Sons, 1986.

[93] Jack W. Macki, Paolo Nistri, and Pietro Zecca. Mathematical models for hysteresis. *SIAM Review*, 35(1):94–123, March 1993.

[94] Akhil J. Madhani, Günter Niemeyer, and J. Kenneth Salisbury. The Black Falcon: A teleoperated surgical instrument for minimally invasive surgery. In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 936–944, October 1998.

[95] Rolf Mamen. Applying space technologies for human benefit; the Canadian experience and global trends. In *Proc. IEEE International Conference on Recent Advances in Space Technologies*, pages 1–8, November 2003.

[96] William R. Mark, Scott C. Randolph, Mark Finch, James M. Van Verth, and Russel M. Taylor II. Adding force feedback to graphics systems: Issues and solutions. In *Proc. SIGGRAPH*, pages 447–452, August 1996.

[97] Thomas Massie. A tangible goal for 3D modeling. *IEEE Computer Graphics and Applications*, 18(3):62–65, May/June 1998.

[98] Thomas H. Massie and J. Kenneth Salisbury, Jr. The PHANToM haptic interface: A device for probing virtual objects. In *Proc. ASME Dynamic Systems and Control, International Mechanical Engineering Congress and Exposition*, volume 55(1), pages 295–301, November 1994.

[99] Michael J. Massimino. *Sensory Substitution for Force Feedback in Space Teleoperation.* Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, 1992.

[100] Michael J. Massimino and Thomas B. Sheridan. Teleoperator performance with varying force and visual feedback. *Human Factors*, 36(1):145–157, March 1994.

[101] Maxon Precision Motors, Inc., http://www.maxonmotorusa.com.

[102] P. R. McAree and R. W. Daniel. Stabilizing impacts in force reflecting teleoperation using distance-to-impact estimates. *International Journal of Robotics Research*, 19(4):349–364, April 2000.

[103] Dwight Meglan. Making surgical simulation real. *Computer Graphics*, 30(4):37–39, November 1996.

[104] Margaret Diane Rezvan Minsky. *Computational Haptics: The Sandpaper System for Synthesizing Texture for a Force-Feedback Display.* Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, 1995.

[105] Manuel Moreyra and Blake Hannaford. A practical measure of dynamic response of haptic devices. In *Proc. IEEE International Conference on Robotics and Automation*, volume 1, pages 369–374, May 1998.

[106] S. Munir, L. Tognetti, and Wayne J. Book. Experimental evaluation of a new braking system for use in passive haptic displays. In *Proc. American Control Conference*, volume 6, pages 4456–4460, June 1999.

[107] Günter Niemeyer. *Using Wave Variables in Time Delayed Force Reflecting Teleoperation*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, September 1996.

[108] Günter Niemeyer, Katherine J. Kuchenbecker, Raymond Bonneau, Probal Mitra, Andrew Reid, Jonathan Fiene, and Grant Weldon. THUMP: An immersive haptic console for surgical simulation and training. In *Proc. Medicine Meets Virtual Reality*, pages 272–274, January 2004.

[109] Günter Niemeyer and Jean-Jacques E. Slotine. Stable adaptive teleoperation. *IEEE Journal of Oceanic Engineering*, 16(1):152–162, January 1991.

[110] Günter Niemeyer and Jean-Jacques E. Slotine. Designing force reflecting teleoperators with large time delays to appear as virtual tools. In *Proc. IEEE International Conference on Robotics and Automation*, volume 3, pages 2212–2218, April 1997.

[111] Günter Niemeyer and Jean-Jacques E. Slotine. Towards force-reflecting teleoperation over the Internet. In *Proc. IEEE International Conference on Robotics and Automation*, volume 3, pages 1909–1915, May 1998.

[112] Günter Niemeyer and Jean-Jacques E. Slotine. Telemanipulation with time delays. *International Journal of Robotics Research*, 23(9):873–890, September 2004.

[113] Allison M. Okamura, Mark R. Cutkosky, and Jack Tigh Dennerlein. Reality-based model for vibration feedback in virtual environments. *IEEE/ASME Transactions on Mechatronics*, 6(3):245–252, September 2001.

[114] Allison M. Okamura, Jack T. Dennerlein, and Robert D. Howe. Vibration feedback models for virtual environments. In *Proc. IEEE International Conference on Robotics and Automation*, volume 3, pages 674–679, May 1998.

[115] Allison M. Okamura, M. W. Hage, Mark R. Cutkosky, and Jack T. Dennerlein. Improving reality-based models for vibration feedback. In *Proc. ASME Dynamic Systems and Controls Division, International Mechanical Engineering Congress and Exposition*, volume 2, pages 1117–1124, November 2000.

[116] Allison M. Okamura, Christopher Richard, and Mark R. Cutkosky. Feeling is believing: Using a force-feedback joystick to teach dynamic systems. *ASEE Journal of Engineering Education*, 91(3):345–349, July 2002.

[117] Marcia K. O'Malley and Michael Goldfarb. On the ability of humans to haptically identify and discriminate real and simulated objects. *Presence: Teleoperators and Virtual Environments*, 14(3):366–376, June 2005.

[118] Rieko Osu and Hiroaki Gomi. Multijoint muscle regulation mechanisms examined by measured human arm stiffness and EMG signals. *Journal of Neurophysiology*, 81(4):1458–1468, April 1999.

[119] Randy Pausch, Thomas Crea, and Matthew Conway. A literature survey for virtual environments: military flight simulator visual systems and simulator sickeness. *Presence: Teleoperators and Virtual Environments*, 1(3):344–363, Summer 1992.

[120] Dianne T. V. Pawluk and Robert D. Howe. Dynamic lumped element response of the human fingerpad. *Journal of Biomechanical Engineering*, 121(2):178–183, April 1999.

[121] Platon A. Prokopiou, William S. Harwin, and Spyros G. Tzafestas. Enhancement of a telemanipulator design with a human arm model. In Spyros G. Tzafestas and G. Schmidt, editors, *Progress in System and Robot Analysis and Control Design*, volume 243 of *Lecture Notes in Control and Information Sciences*, pages 445–456. Springer-Verlag, 1999.

[122] Platon A. Prokopiou, Spyros G. Tzafestas, and William S. Harwin. Towards variable-time-delays-robust telemanipulation through master state prediction. In *Proc. IEEE/ASME International Conference on Advanced Intelligent Mechanisms*, pages 305–310, September 1999.

[123] G. Jagannath Raju, George C. Verghese, and Thomas B. Sheridan. Deisgn issues in 2-port network models of bilateral remote manipulation. In *Proc. IEEE International Conference on Robotics and Automation*, volume 3, pages 1316–1321, May 1989.

[124] Denis Rancourt and Neville Hogan. Stability in force-production tasks. *Journal of Motor Behavior*, 33(2):193–204, June 2001.

[125] Jens Rassweiler, Jochen Binder, and Thomas Frede. Robotic and telesurgery: Will they change our future? *Current Opinion in Urology*, 11(3):309–320, May 2001.

[126] Howard Rheingold. *Virtual Reality*. Simon and Schuster, New York, 1991.

[127] Dipartimento di Ingegneria Aerospaziale, Politecnico di Milano, RTAI: Realtime Application Interface for Linux, http://www.aero.polimi.it/∼rtai/.

[128] Septimiu E. Salcudean and Tim D. Vlaar. On the emulation of stiff walls and static friction with a magnetically levitated input-output device. *ASME Journal of Dynamics, Measurement and Control*, 119:127–132, March 1997.

[129] J. Kenneth Salisbury, Jr. The heart of microsurgery. *Mechanical Engineering Magazine*, 120(12):47–51, December 1998.

[130] J. Kenneth Salisbury, Jr., William Townsend, Brian Eberman, and David DiPietro. Preliminary design of a whole-arm manipulator system (WAMS). In *Proc. IEEE International Conference on Robotics and Automation*, volume 1, pages 254–260, April 1988.

[131] Sava Industries, Inc., http://www.savacable.com.

[132] SensAble Technologies, Inc., http://www.sensable.com.

[133] Thomas B. Sheridan. Musings on telepresence and virtual presence. *Presence: Teleoperators and Virtual Environments*, 1(1):120–126, Winter 1992.

[134] Alana Sherman, Murat Cenk Çavuşoğlu, and Frank Tendick. Comparison of teleoperator control architectures for palpation task. In *Proc. ASME Dynamic Systems and Control Division, International Mechanical Engineering Congress and Exposition*, volume 2, pages 1261–1268, November 2000.

[135] Carl E. Sherrick and Roger W. Cholewiak. Cutaneous sensitivity. In Kenneth R. Boff, Lloyd Kaufman, and James P. Thomas, editors, *Handbook of Perception and Human Performance*, volume I: Sensory Processes and Perception, chapter 12. John Wiley and Sons, 1986.

[136] Karun B. Shimoga. Finger force and touch feedback issues in dexterous telemanipulation. In *Proc. Fourth Annual Conference on Intelligent Robotic Systems for Space Exploration*, pages 159–178, September 1992.

[137] Mohammad Reza Sirouspour and Septimiu E. Salcudean. Suppressing operator-induced oscillations in manual control systems with movable bases. *IEEE Transactions on Control Systems Technology*, 11(4):448–458, July 2003.

[138] Metin Sitti and Hideki Hashimoto. Teleoperated touch feedback from the surfaces at the nanoscale: Modeling and experiments. *IEEE/ASME Transactions on Mechatronics*, 8(2):287–298, June 2003.

[139] Mandayam A. Srinivasan, Cagatay Basdogan, and Chih-Hao Ho. Haptic interactions in virtual worlds: Progress and prospects. In *Proc. International Conference on Smart Materials, Structures, and Systems*, July 1999.

[140] Gyung Tak Sung and Inderbir S. Gill. Robotic laparoscopic surgery: a comparison of the da Vinci and Zeus systems. *Urology*, 58(6):893–898, December 2001.

[141] Neal A. Tanner and Günter Niemeyer. Online tuning of wave impedance in telerobotics. In *Proc. IEEE Conference on Robotics, Automation, and Mechatronics*, pages 7–12, December 2004.

[142] Neal A. Tanner and Günter Niemeyer. Improving perception in time delayed teleoperation. In *Proc. IEEE International Conference on Robotics and Automation*, pages 356–361, April 2005.

[143] Russell Taylor, Pat Jensen, Louis Whitcomb, Aaron Barnes, Rajesh Kumar, Dan Stoianovici, Puneet Gupta, ZhengXian Wang, Eugene deJuan, and Louis Kavoussi. A steady-hand robotic system for microsurgical augmentation. *International Journal of Robotics Research*, 18(12):1201–1210, December 1999.

[144] Frank Tendick and Murat Cenk Çavuşoğlu. Human-machine interfaces for minimally invasive surgery. In *Proc. IEEE Engineering in Medicine and Biology Society*, pages 2771–2776, November 1997.

[145] Frank Tendick, Michael Downes, Tolga Goktekin, Murat Cenk Çavuşoğlu, David Feygin, Xunlei Wu, Roy Eyal, Mary Hegarty, and Lawrence W. Way. A virtual environment testbed for training laparoscopic surgical skills. *Presence: Teleoperators and Virtual Environments*, 9(3):236–255, June 2000.

[146] William T. Townsend and J. Kenneth Salisbury, Jr. Mechanical bandwidth as a guideline to high-performance manipulator design. In *Proc. IEEE International Conference on Robotics and Automation*, volume 3, pages 1390–1395, May 1989.

[147] Toshio Tsuji, Kazuhiro Goto, Masamitsu Moritani, Makoto Kaneko, and Pietro Morasso. Spatial characteristics of human hand impedance in multi-joint arm movements. In *Proc. IEEE International Conference on Intelligent Robots and Systems*, volume 1, pages 423–430, September 1994.

[148] Alan M. Turing. Computing machinery and intelligence. *Mind: A Quarterly Review of Psychology and Philosophy*, 59(236):433–460, October 1950.

[149] Jean Vertut, J. Charles, Philippe Coiffet, and M. Petit. Advance of the new MA 23 force reflecting manipulator system. In *Proc. Second CISM-IFToMM Symposium on Theory and Practice of Robots and Manipulators*, volume 3, pages 307–317, September 1976.

[150] Holly S. Vitense, Julie A. Jacko, and V. Kathlene Emery. Multimodal feedback: an assessment of performance and mental workload. *Ergonomics*, 46(1–3):68–87, January 2003.

[151] Christopher R. Wagner, Nicholas Stylopoulus, and Robert D. Howe. The role of force feedback in surgery: Analysis of blunt dissection. In *Proc. IEEE Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 68–74, March 2002.

[152] Parris S. Wellman and Robert D. Howe. Towards realistic display in virtual environments. In *Proc. ASME Dynamic Systems and Control Division*, volume 57, pages 713–718, 1995.

[153] Aaron M. West and Mark R. Cutkosky. Detection of real and virtual fine surface features with a haptic interface and stylus. In *Proc. ASME Dynamic Systems and Control Division*, volume 61, pages 159–166, November 1997.

[154] G. Westling and Roland S. Johansson. Responses in glabrous skin mechanoreceptors during precision grip in humans. *Experimental Brain Research*, 66(1):128–140, March 1987.

[155] Howard Wisniowski. Analog devices puts micromachines in motion. Technical report, Analog Devices, Inc., http://www.analog.com, 1996.

[156] Yasuyoshi Yokokohji and Tsuneo Yoshikawa. Bilateral control of master-slave manipulators for ideal kinesthetic coupling - formulation and experiments. *IEEE Transactions on Robotics and Automation*, 10(5):605–620, October 1994.

[157] Tsuneo Yoshikawa and Akihiro Nagura. A touch/force display system for haptic interface. *Presence: Teleoperators and Virtual Environments*, 10(2):225–235, April 2001.

[158] Craig Zilles and J. Kenneth Salisbury, Jr. A constraint based god-object method for haptic display. In *Proc. International Conference on Intelligent Robots and Systems*, 1995.