

THE EFFECTS OF DISPLACEMENT QUANTIZATION AND
ZERO-ORDER HOLD ON THE
LIMIT CYCLE BEHAVIOR OF HAPTIC KNOBS

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

Christopher J. Hasser

December 2001

© Copyright 2002 by Christopher J. Hasser
All Rights Reserved

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, Principal Adviser

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.

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. Christian Gerdes

Approved for the University Committee on Graduate Studies:

Abstract

Haptic feedback devices measure human motion using displacement sensors such as optical encoders, and use actuators to apply computer-programmable forces to the user. Haptic devices act as both displays and input devices, and haptic system control loops include human operator dynamics. Limit cycle oscillations upon contact with a virtual barrier are a common problem in haptic devices, contributed to by factors such as discrete sampling effects and displacement quantization. Several investigators have examined discrete sampling effects, but a knowledge gap exists regarding the interaction of these two factors, particularly for coarse-resolution displacement sensors common in commercial haptic systems. As a prerequisite to simulations and theoretical analysis of a haptic knob in contact with a virtual barrier (unilateral spring) for various combinations of sample rate and displacement quantization, the present dissertation includes system identification experiments for the human hand grasping a haptic knob. A second-order rotational model fits the data for mild to moderate grasp forces, and a fourth-order model that treats the fingerpad and finger separately is suggested for strong grip forces. Simulations reveal that increasing displacement quantization increases limit cycle oscillation magnitude but has relatively little effect on limit cycle frequency. Increasing sample rate decreases oscillation magnitude but can easily increase limit cycle frequencies into ranges more easily perceptible by the sense of touch or hearing. Empirical data gathered with a haptic knob and human user support the simulation

results. Analysis shows that displacement quantization exacerbates the well-known energy-instilling effects of the zero-order hold inherent in a computer-driven haptic system, especially upon entry and exit from the virtual barrier. The product of encoder resolution and sample period is shown to be a good predictor of the sensitivity of limit cycle oscillation magnitude to sample rate and displacement quantization. Describing function analysis is used to provide accurate predictions of limit cycle magnitude and frequency. Methods for mitigating limit cycle oscillations in systems with coarse displacement sensing, and other implications for design, are discussed, emphasizing the desirability of reducing limit cycle oscillation magnitude without increasing oscillation frequency into ranges of heightened sensitivity for the senses of touch or hearing.

Acknowledgements

Mark Cutkosky, my advisor, presented me with the wonderful opportunity to continue my studies at Stanford. I have been grateful for his advice and counsel, as well as his flexibility as I have juggled many responsibilities. Chris Gerdes offered insights and encouragement that were particularly valuable during the system identification portion of this work. His advisees are lucky men and women. Ken Salisbury has made his warm and intelligent presence felt in personal ways over the years, helped directly with enthusiasm and advice during the writing of this dissertation, and provided indirect support through his contributions to the field of haptics. Gunter Niemeyer acted as an unofficial committee member. I appreciated his help and accessibility during the crucial last weeks of my dissertation work.

The work of two individuals stands out among all others in creating fertile territory to develop the present dissertation. Aram Hajian's dissertation performed under the supervision of Professor Rob Howe set an inspirational example, and provided valuable precedent for the system ID experiments performed here. Professor Howe has been a source of sage advice for the past decade. Brent Gillespie's dissertation performed under the supervision of Mark Cutkosky was exciting for its relevance and also of a quality that I found inspirational. Brent generously gave me Matlab code that he used for his "energy

leak” modeling work. This code provided a kernel from which to grow the code developed for the simulation work in the latter half of this dissertation.

I owe a great debt of gratitude to my mentors at the United States Air Force Armstrong Laboratory: Maj. Ron Julian, Lt. Col. Peter Lurker, Dr. Tom Moore, Dr. Charles Nixon, and Lt. Col. (Dr.) Paul Whalen. They took a gamble on an Air Force cadet from Cornell, and created a haptics researcher. I owe a similar debt to the haptics research community. It is a “community” in the truest sense; one that I found stimulating and supportive. Dov Adelstein, Ed Colgate, Jack Dennerlein, Blake Hannaford, Vincent Hayward, Susan Lederman, Karon MacLean, Tim Salcudean, Mandayam Srinivasan, Frank Tendick, Jan Weisenberger, their students, and others have been generous with their help for many years. Malcom Daniels, my master’s thesis advisor at the University of Dayton, came through on a promise to offer a controls curriculum and played a pivotal role in my development.

The folks at Immersion Corporation have set the highest standard of technical excellence, have provided a context to motivate this work, and have been personal friends. Louis Rosenberg first showed me his drive and strategic mind as a graduate fellow at the Armstrong Laboratory in Ohio, and extended the offer for me to join him on his Great California Startup Adventure. Beginning with a rental truck, a flat tire and a temporarily lost cat on the journey West, it has been a wild ride. It was a long time ago that we visited Young’s Dairy in Yellow Springs Ohio with our girlfriends Sheila and Sherstin (now our wives and the mothers of our children).

Bruce Schena has been an excellent supervisor and mentor. I will measure all engineers against the example he has set, and appreciate our shared values regarding the

mission of a Research & Development program. His support has played a major role in making the attainment of this goal possible. His design of the Haptic Reference (HREF) hardware was crucial to the performance of the system identification work in this dissertation. Danny Grant and Oliver Astley have been valued R&D compatriots and virtual labmates. They also contributed to the HREF system. Mike Brown, Alex Goldenberg, Alex Jasso, Jeff Mallet, Heinz Saurer, Greg Springer, Kollin Tierling, and Steve Vassallo provided assistance with equipment tutorials, insights, machining help, and moral support. I could not have done this without Jonathan Chu Yang's reliable network and data storage. Adam Braun, Dean Chang, Mike Levin, Ken Martin, Erik Shahoian, and Evan Wies set the bar very high, and it has been an honor to work with them. Craig Factor, David Hague, and Tim Lacey will also remain an unforgettable part of this journey.

Fellow students in the Touch Group, Jorge Cham, Mike Costa, Wes Griffin, Allison Okamura, Sile O'Modhrain, Chris Richard, Neils Smaby, and Mike Turner set a good example, provided technical and moral support and encouragement, and proved to me that it is possible to graduate. I wish I had been able to spend more time with them.

Engineers at MicroE Systems, Bill Marshall, Matt Giza, Michael Foley, and Paul Tullar, provided crucial and patient help as I got their new Mercury sensor up and running and dealt with various issues. Especially in a dissertation on low-resolution displacement sensing, their ultra-high resolution encoder was a pleasure to use.

My parents deserve much credit for this achievement. They fed my curiosity from an early age, and gave me countless opportunities to grow. I well remember early trips to the library with my mother. My father's story, oft repeated, of his graduate assistant job

at the University of Missouri filing the records of students who had completed all requirements for a degree except for the thesis, must be mentioned here. I shall consider myself forever immune to further repetition of that story. The Boy Scouts of America were another early influence critical to this achievement. Scouting taught me to climb a mountain, physical or mental, one step at a time, resolutely resisting the infinite opportunities to quit, in pursuit of singular success.

Finally, my wonderful family has supported me through this work, and tolerated my absence. I appreciate my wife's wisdom, insight, patience, and support. This achievement is one of many things in my life that would not have happened without her. The sweetest reward for completing this dissertation will be the ability to spend more time with my family.

Table of Contents

Abstract	v
Acknowledgements	vii
Table of Contents	xi
List of Figures	xiii
Chapter 1 : Introduction	1
1.1 Motivation	3
1.2 Dissertation Contributions	6
1.3 Dissertation Overview	7
Chapter 2 : Human System Identification Methods and Literature	9
2.1 System Identification Input Waveforms	13
2.2 Previous Studies of Finger Dynamics	15
Chapter 3 : Apparatus	19
3.1 Hardware Setup	19
3.2 Software Architecture	21
Chapter 4 : System Identification for a Fingertip Pinch Grasp	23
4.1 Introduction	23
4.2 Methods	23
4.2.1 Apparatus and Procedure	23
4.2.2 Determination of the Relevant Displacement Range	25
4.2.3 System Identification Technique	26
4.3 Results	28
4.4 Discussion	34
4.4.1 Suggested 4 th Order Model	39
4.5 Conclusions	45
Chapter 5 : System Identification for More Grasps and Conditions	46
5.1 Introduction	46
5.2 Methods	46
5.3 Results and Discussion	48
5.4 Conclusions	53
Chapter 6 : Simulation of Virtual Barrier Contact for Various Displacement Resolutions and Sample Rates	56
6.1 Introduction	56
6.2 Previous Work	57
6.3 Does Position Quantization Possess Inherent Energy Leaks?	65
6.4 “Energy Leak” Analysis for a Haptic Knob	70
6.4.1 Initial Simulations with Inertia, No Damping or Hand Stiffness (The “Bouncing Ball” Model)	71

6.4.2 With Hand Dynamics Added	74
Chapter 7 : Describing Function Analysis	84
7.1 Introduction	84
7.2 Application of Describing Function Analysis to Haptic Barrier Limit Cycles	85
7.3 Results	91
7.4 Discussion	97
Chapter 8 : Hardware Testing	99
8.1 Introduction	99
8.2 Methods	99
8.3 Results	100
8.4 Discussion	106
Chapter 9 : Effects of Sample Rate and Displacement Quantization on the Amplitude and Frequency of Limit Cycle Oscillations	108
9.1 Why Does Displacement Quantization Worsen ZOH-Induced Energy Leaks? ...	108
9.2 Why Do the Bouncing Ball Simulations Show Less Sensitivity to Encoder Resolution Than the Hand-on-Knob Simulations?	113
9.3 Prediction of Limit Cycle Oscillation Magnitude	114
9.4 Frequency Behavior	118
9.5 The Role of Friction	119
Chapter 10 : Design Implications	120
10.1 Introduction	120
10.2 Design Variables Affecting Limit Cycle Behavior	121
10.3 Touch and Hearing	123
10.4 Design Implications	125
10.5 An Example to Discourage High Sample Rates and Low Displacement Resolutions	126
10.6 Liabilities of Sampling Too Quickly	129
10.7 A Suggested Approach	130
10.8 Other Design Implications	133
Chapter 11 : Limit Cycle Mitigation Approaches	134
11.1 Typical Oscillations	137
11.2 Impulse Damping	139
11.3 Phase Estimation Damping	143
11.3.1 Initial Phase Estimation Damping Results	145
11.4 Implications for Multi-Degree-of-Freedom Systems	147
Chapter 12 : Conclusions	149
Appendix A : Motor Calibration	152
A.1 Torque Constant and Dynamic Friction	152
A.2 Transient Response	155
Appendix B : Grip Force Load Cell Calibration	157
Appendix C : Apparatus Dynamics	160
Appendix D : Raw Data for Empirical Trials	163
References	171

List of Figures

Figure 1: Haptic Reference (HREF) knob apparatus	20
Figure 2: Subject's hand grasping HREF knob in a pinch grasp.....	25
Figure 3: Oscillations Captured from Haptic Knob with Detents.....	26
Figure 5: Thumbnail summary plots of displacement trajectories for all 54 trials for a particular subject; 9 trials per grip force condition, with increasing grip force from left-right and top-bottom.....	30
Figure 6: Commanded torque (upper left), acceleration (upper right), velocity (lower left), and displacement vectors (lower right) for a typical trial with a grip force of 2.4 N.....	30
Figure 7: Commanded torque ($T_{q_{cmd}}$), total estimated torque ($T_{q_{calc}}$), and estimated torque contributions from moment of inertia ($T_{q_{acc}}$), damping ($T_{q_{vel}}$), and displacement ($T_{q_{disp}}$) [plot layout adapted from Hajian, 1997].	31
Figure 8: Measured step responses for a trial at six different grip forces (solid lines), simulated trajectories based on model estimates derived from the same data (dotted lines), and simulated trajectories based on models derived using data from a different trial (dot-dashed lines).....	31
Figure 9: Moment of Inertia, Damping, Stiffness, and Damping Ratio for One Subject	32
Figure 10: Moment of Inertia, Damping, Stiffness, and Damping Ratio for All Subjects	32
Figure 11: Moment of Inertia, Damping, Stiffness, and Damping Ratio Averaged Across All Subjects	33
Figure 12: Comparison of Normalized Results to Those of Hajian's Finger Extension and Abduction Experiments.....	38
Figure 13: Suggested fourth order model, with finger dynamics, fingerpad dynamics, and actuator dynamics separately represented.	40
Figure 14: Fourth order mechanical network.....	40
Figure 15: Displacement Responses of Estimated 4th and 2nd-Order Systems Compared to Measured Response.....	43
Figure 16: Acceleration Responses of Estimated 4th and 2nd-Order Systems Compared to Measured Response.....	44
Figure 17: Photographs of four knob grasps. (a) key pinch on 17.8 mm knob; (b) simple pinch on 29.9 mm knob; (c) key pinch on 29.9 mm knob; (d) three-fingered pinch on 29.9 mm knob.....	47
Figure 18: Comparison of Experiments for Different Grasps on the 17.8 mm Knob Experiment 1, Simple Pinch Grasp (solid line) Experiment 2, Key Pinch Grasp (dotted line)	49

Figure 19: Comparison of Experiments for Different Grasps on the 29.9 mm Knob Experiment 3, Simple Pinch Grasp (solid line) Experiment 4, Key Pinch Grasp (dotted line) Experiment 5, Three-Finger Grasp (dot-dashed line)	51
Figure 20: Comparison of Experiments for All Five Experiments Experiment 1, Simple Pinch Grasp (solid line, with '●' data points) Experiment 2, Key Pinch Grasp (dotted line, with '●' data points) Experiment 3, Simple Pinch Grasp (solid line, with 'O' data points) Experiment 4, Key Pinch Grasp (dotted line, with 'O' data points) Experiment 5, Three-Finger Grasp (dot-dashed, with 'O' data points)	52
Figure 21: Comparison of Experiments for All Five Experiments, converted to NORMALIZED TRANSLATIONAL MODELS Experiment 1, Simple Pinch Grasp (solid line, with '●' data points) Experiment 2, Key Pinch Grasp (dotted line, with '●' data points) Experiment 3, Simple Pinch Grasp (solid line, with 'O' data points) Experiment 4, Key Pinch Grasp (dotted line, with 'O' data points) Experiment 5, Three-Finger Grasp (dot-dashed, with 'O' data points)	53
Figure 22: Plot of modeled manipulandum position and control effort (from Gillespie and Cutkosky, 1996).	60
Figure 23: Continuous-Time Simulation with Encoder Displacement Quantization	66
Figure 24: Continous-Time Simulations of a Very Stiff Wall with Displacement Quantization	68
Figure 25: Closeup of a Bounce During Continous-Time Simulation of a Very Stiff Wall with Displacement Quantization	68
Figure 26: Continuous-Time Simulation with Encoder Displacement Quantization and Physical Damping	69
Figure 27: Gillespie's Model of a Finger/Manipulandum Contacting a Virtual Wall (from Gillespie, 1996)	70
Figure 28: "Bouncing Ball" Simulation of Manipulandum Contacting Virtual Wall with Time Discretization and Displacement Quantization, but no Physical Damping. Sample Rate = 1 kHz; Encoder Resolution = 4096 counts/revolution.....	72
Figure 29: Quadratic Energy Growth Rates as a Function of Sampling Period and Encoder Resolution	73
Figure 30: Quadratic Energy Growth Rates as a Function of Sampling Period and Encoder Resolution (Log Magnitude for Growth Rate).....	74
Figure 31: Simulated Oscillations of a Lightly-Grasped Haptic Knob in Contact with a Stiff Virtual Barrier	76
Figure 32: Simulation of Hand Lightly Pressing Knob Against Stiff Virtual Wall, with Lines Fitted to Steady State Peaks and Troughs to Measure Limit Cycle Magnitude (2000 Hz, 8192 encoder counts/revolution).....	77
Figure 33: Peak-to-Peak Oscillation Magnitude for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Period and Displacement Quantization	79
Figure 34: Peak-to-Peak Oscillation Magnitude for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Rate (log10) and Displacement Counts/Revolution (log2)	79

Figure 35: Peak-to-Peak Oscillation Magnitude (log10) for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Rate (log10) and Displacement Counts/Revolution (log2).....	80
Figure 36: Peak-to-Peak Oscillation Magnitude, Expressed in Units of Encoder Counts	80
Figure 37: Saturated Mesh Showing Cases with Oscillation Magnitudes Greater than +/- 1 Encoder Count as "1" and Cases with Oscillations Less Than +/- 1 Encoder Count as "-1"	81
Figure 38: Oscillation Frequency for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Rate (log10) and Displacement Counts/Revolution (log2).....	82
Figure 39: Average Limit Cycle Frequency as a Function of Servo Rate	82
Figure 40: Average Limit Cycle Frequency as a Function of Displacement Resolution.	83
Figure 41: Block diagram of a closed-loop linear system with linear component, $G(j\omega)$, and nonlinear component, $H(j\omega)$	85
Figure 42: The relay nonlinearity.....	86
Figure 43: Virtual barrier nonlinearity	88
Figure 44: Virtual barrier nonlinearity with bias torque included	89
Figure 45: Nyquist plot of linear system with nominal describing function (dashed line) and delayed describing function (angled line with circle denoting intercept)	90
Figure 46: Nyquist plot (curve) with angled lines plotting describing functions for various ZOH delays, with circles indicating intersections.....	92
Figure 47: Log magnitude of limit cycles predicted by describing function analysis	93
Figure 48: Log magnitude of limit cycles predicted by describing function analysis (left) compared to Log magnitude of limit cycles predicted by simulations (right)	93
Figure 49: Difference between DFA and simulation magnitudes as a percentage of simulation magnitudes.....	94
Figure 50: Frequencies of limit cycles predicted by describing function analysis	96
Figure 51: Frequencies of limit cycles predicted by describing function analysis (left) compared to frequencies of limit cycles predicted by simulations (right)	96
Figure 52: Difference between DFA and simulation frequencies as a percentage of simulation frequencies.....	97
Figure 53: Grip Force for Each of the Sixteen Trials.....	101
Figure 54: Thumbnail Sketches of Raw Displacement Data	102
Figure 55: RMS Magnitude of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution	103
Figure 56: Log10 peak-to-peak Magnitude of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution.....	103
Figure 57: A comparison of simulation to hardware results for oscillation magnitude as a function of sample rate and displacement resolution.....	104
Figure 58: Difference between hardware testing and simulation magnitudes as a percentage of simulation magnitudes.....	105
Figure 59: Principal Frequency of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution	106
Figure 60: Illustration of Barrier Penetration and Resultant Torque Outputs for a Traditional ZOH System and a ZOH System with Displacement Quantization.....	110

Figure 61: Quadratic Energy Growth Rates for the Bouncing Ball Case (left) and Peak-to-Peak Magnitudes for the Hand-on-Knob Case (right)	114
Figure 62: Simulation Results for Peak-to-Peak Oscillation Magnitude	116
Figure 63: Predictions for Peak-to-Peak Oscillation Magnitude	116
Figure 64: Log10 RMS Magnitude of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution	117
Figure 65: Log10 RMS Magnitude of Predictor Function Over Same Range of Parameters as Hardware Trials.....	117
Figure 66: Average Limit Cycle Frequency as a Function of Servo Rate	118
Figure 67: Equal Loudness Contours for the Sense of Hearing, Based On the Work of Fletcher and Munson (image: D. Cross, in Gotfrit 2001, by permission).....	125
Figure 68: Simulation Results Showing Oscillation Magnitude for Current Design (left corner, 2650 count/rev, 1 kHz) and Test Case (right corner, 360 counts/rev, 20 kHz)	127
Figure 69: Limit Cycle Oscillations for Virtual Barrier at 25 kHz Sample Rate with 360 counts/revolution Encoder, Recorded from the HREF Hardware.....	129
Figure 70: Notional Optimization Surface.....	132
Figure 71: Time Response of Oscillations with 1 kHz Sample Rate and 2650 counts/rev	138
Figure 72: Phase Plane Plot of Oscillations with 1 kHz Sample Rate and 2650 counts/rev	138
Figure 73: Simulated Limit Cycle Oscillations at 1 kHz Sample Rate, with 2650 counts/revolution Encoder, with Impulse Damping Applied.....	141
Figure 74: Phase Plane Plot for Impulse Damping Example.....	142
Figure 75: Time Response Plot for Initial Attempt at Phase Estimation Damping	145
Figure 76: Phase Plane Plot for Initial Attempt at Phase Estimation Damping.....	146
Figure 77: Motor Torque and Dynamic Friction as a Function of Drive Current for Small Torques (less than 2 mNm)	154
Figure 78: Motor Torque and Dynamic Friction as a Function of Drive Current for the Relevant Operating Range (0-30 mNm)	154
Figure 79: Transient Response of the Amplifier and Motor in HREF.....	156
Figure 80: (a) Knob with calibration beam; (b) Closeup of knob.....	157
Figure 81: Grip Force Load Cell Calibration Plot (with calibration masses doubled to account for lever arm in experimental setup).....	159
Figure 82: Raw Data for Trial at 455 Hz and 256 counts/revolution.....	163
Figure 83: Raw Data for Trial at 1000 Hz and 256 counts/revolution.....	163
Figure 84: Raw Data for Trial at 2 kHz and 256 counts/revolution.....	164
Figure 85: Raw Data for Trial at 4 kHz and 256 counts/revolution.....	164
Figure 86: Raw Data for Trial at 455 Hz and 512 counts/revolution.....	165
Figure 87: Raw Data for Trial at 1 kHz and 512 counts/revolution.....	165
Figure 88: Raw Data for Trial at 2 kHz and 512 counts/revolution.....	166
Figure 89: Raw Data for Trial at 5 kHz and 512 counts/revolution.....	166
Figure 90: Raw Data for Trial at 455 Hz and 1024 counts/revolution.....	167
Figure 91: Raw Data for Trial at 1 kHz and 1024 counts/revolution.....	167
Figure 92: Raw Data for Trial at 2 kHz and 1024 counts/revolution.....	168
Figure 93: Raw Data for Trial at 5 kHz and 1024 counts/revolution.....	168

Figure 94: Raw Data for Trial at 455 Hz and 2048 counts/revolution.....	169
Figure 95: Raw Data for Trial at 1 kHz and 2048 counts/revolution.....	169
Figure 96: Raw Data for Trial at 2 kHz and 2048 counts/revolution.....	170
Figure 97: Raw Data for Trial at 5 kHz and 2048 counts/revolution.....	170

Chapter 1 : Introduction

The word *haptic* comes from the Greek *haptesthai*, to touch. In current use, it means “of or related to the sense of touch.” The terms *haptics* and *haptic feedback* have come to mean the field of research and products intended to provide computer-driven displays to the sense of touch through the use of actuated devices such as exoskeletons, joysticks, steering wheels, and knobs. Many of the developments in haptics trace their roots to the field of force-feedback teleoperation, pioneered by Raymond Goertz at the Argonne National Laboratory shortly after World War II, in order to handle “hot” nuclear materials with remote robots while allowing operators to feel the forces experienced by the robots. In 1965, Ivan Sutherland predicted the emergence of computer-driven displays for the sense of touch:

I want to describe for you a kinesthetic display. The force required to move a joystick could be computer controlled, just as the actuation force on the controls of a Link Trainer are changed to give the feel of a real airplane... By use of such an input/output device, we can add a force display to our sight and sound capability.

(Sutherland, 1965)

In the late 1980’s and early 1990’s, the field was dominated by complex devices like full-arm exoskeletons, and high-performance devices such as machined aluminum joystick mechanisms with cable transmissions and rare-earth magnet DC motors. Cost

was generally not a factor in design. Research and high-end simulation applications such as for surgery education still have only modest cost pressures on design. While desired performance increases have driven some component technology development, particularly in the area of actuators, cost concerns have not motivated component technology development for these applications.

In the late 1990's several companies began to think about commercializing haptic devices for mass-market applications. Force-feedback joysticks, steering wheels, and mice for video games and other computer applications began to appear on the market. Cost began to battle performance in the minds of commercial haptic system designers. Proponents of commercial haptics found themselves attempting to introduce an enhancement to commodity consumer products with intense pressures to minimize cost.

The evolution of haptic systems generally followed a course from many degrees of freedom (e.g., Goertz' manipulators) to fewer degrees of freedom (e.g., joysticks and steering wheels). The cost of haptic devices scales nonlinearly with the number of degrees of freedom. Low degree-of-freedom devices have the potential to reach more people in more applications at a lower cost. The extreme example of this is a single degree-of-freedom device such as a haptic knob for embedded applications. Good and Munson (1993) did some early work on a single degree-of-freedom haptic torque device designed to display molecular interaction forces. Schuler (1993, 1995, 1996) developed a haptic feedback knob with programmable detents for a video editing console. Immersion Corporation first began to commercialize haptic knobs through licensing to automobile companies like BMW (Ashley, 2001). The iDrive haptic knob in 7-series 2002 BMW sedans works with a dashboard-mounted graphical user interface to control over 700

secondary vehicle functions for subsystems such as heating, cooling, stereo, navigation, and seat controls.

1.1 Motivation

The present dissertation grew out of an interest in using research to predict and improve the performance of haptic devices designed with stringent cost constraints. Such constraints often encourage the selection of less expensive, lower performance components; however, in many cases performance requirements cannot be relaxed. This is especially true for haptic knobs in high-end automobiles. Sensors and actuators are the two cornerstone technologies for haptics. Both drive the performance and cost of haptic systems. Sensors are particularly interesting from a controls perspective. Digital rotation sensors such as incremental optical encoders are ubiquitous in haptic applications, and increasing resolution, their key performance metric, directly increases their cost.

Just as those concerned with actuators would love to have a point torque source with zero mass, zero volume, and infinite torque, those concerned with sensors would benefit greatly from a zero cost, zero volume, infinite resolution displacement sensor. Component technology research and development efforts can produce better, less costly, actuators and sensors; however, in accordance with the University of Chicago economists' mantra, there is no such thing as a free lunch. Component technology R&D can carry very high financial costs, which represent large opportunity costs to public funding agencies or for-profit companies with limited development funds.

Questions such as the determination of adequate encoder resolution can have a bearing on the quality of fielded products, and can also help to guide strategic technology

development decisions. Expensive technology development efforts should be motivated by a demonstrated need, and should be informed by requirements that are reasonably known to meet the demonstrated need. Stability-enhancing algorithms can increase the quality of haptic interfaces and may allow quality to remain constant while designers reduce system costs.

This dissertation aims to develop an understanding of the effects of displacement resolution on haptic system performance, with an underlying desire to either discover advanced controls techniques to allow low-resolution systems to perform as well as high-resolution systems, or to develop confidence in the need to invest resources in high-resolution, low-cost sensor technology research. Though displacement sensors influence performance in many ways, this work investigates the ubiquitous case of a haptic manipulandum (a knob, in particular) with displacement quantization and a discrete-time controller contacting a virtual barrier or wall that acts as a unidirectional spring. The displacement quantization and the zero-order hold effects of the discrete controller can produce limit cycle oscillations upon manipulandum contact with the barrier that the user experiences as buzzing or whining that he or she can both feel and hear. This effect, which might be mildly perturbing to the researcher, could completely negate the value of haptics in a luxury automobile. Other factors that can have a bearing on the character of limit cycle oscillations include the quality of the velocity signal, any physical damping that may be present, and the virtual barrier stiffness and damping.

Haptic designers and researchers have long known that the stability and limit cycle behavior of haptic interfaces depends heavily on the biomechanical properties of users' hands. The size of the hand, the strength of the grasp, and other factors can lead to wide

variations in stability. A device that remains stable when held in a moderate to strong grip may begin to chatter when held in a light grip. In order to design haptic devices for inherent stability (e.g., specifying mechanical parameters, determining acceptable sensor or actuator quantization, etc.), or to implement stability-enhancing control schemes, one needs to know the range of mechanical system properties that a user's hand could present to the system.

Accurate models of the human hand in a knob grasp would help towards these ends, and would also provide a valuable simulation tool for testing designs. Design testing using simulation models of hand dynamics has several advantages over iterative testing with hardware:

- Repeatability
- Ability to try many parameters in automated fashion
- Precise control over experimental variables and conditions (e.g., can have hand and manipulandum approach wall with particular velocity, or lean against it with particular force trajectory)
- Ease of obtaining data from a validated model compared to physical hardware
- Ability to analyze hardware designs that have not yet been built

This dissertation includes system identification of the human hand in a knob grasp, development of dynamic models for the hand/knob system, investigation of the effects of displacement quantization and sample rate on limit cycle oscillation magnitude and frequency, discussion of the design implications of these results, and discussion of potential methods for mitigating limit cycle oscillations. The following section enumerates the specific contributions of this dissertation to the body of knowledge in the

field of haptics, and the final section of this chapter gives an overview of the contents of subsequent chapters.

1.2 Dissertation Contributions

The three most significant contributions of this dissertation are:

1. The empirical validation of a second-order rotational model approximation for the dynamics of the human hand grasping a haptic knob, with empirically determined parameters for moment of inertia, damping, and stiffness for several test cases.
2. Discovery and explanation of the effects of displacement quantization on the magnitude and frequency of limit cycle oscillations occurring when a haptic knob contacts a virtual barrier. This work also considers the effects of sample rate (previously considered by other investigators) and the interaction between displacement quantization and sample rate (a new contribution).
3. Development and validation of two tools; the first, a simple approximation of oscillation magnitude that requires knowledge from one empirical case to predict the effect of changing sample rate or displacement quantization, and the second, a successful application of the nonlinear control theory technique of describing function analysis to predict the magnitude and frequency of limit cycles given a dynamic plant model and other system design parameters such as displacement quantization, sample rate, and virtual barrier stiffness.

The dissertation takes these developments and considers their implication for the design of haptic devices. One insight counters the conventional wisdom that faster sample rates are always better by pointing out that faster sample rates have the

disadvantage of increasing limit cycle oscillation frequency, despite their beneficial effect of decreasing oscillation magnitude. The knowledge gained from the present work will help inform traditional haptic device design, but also could provide a useful foundation upon which to base creative approaches to mitigating limit cycle oscillations in systems with relatively low resolution displacement sensors. Several potential solutions for mitigating limit cycles are discussed, and two new techniques are briefly explored, without much success, but with potential for improvement in the future.

1.3 *Dissertation Overview*

The present work develops dynamic models of the human hand grasping a haptic knob, uses simulation, theoretical analysis, and empirical testing to explore the effects of displacement discretization and zero-order hold effects on limit cycle behavior, and discusses the implications for the design of haptic devices and for the guidance of sensor research programs. Chapter 2 covers human system identification methods and prior literature. Chapter 3 describes the apparatus used for the system identification experiments. Chapter 4 presents results for a simple two-finger pinch grasp on a 17.8 mm diameter knob. Chapter 5 presents experiments for different grasp postures and knob diameters. Chapter 6 uses simulation to explore the effects of displacement quantization and sampling rate on limit cycles for a haptic knob in contact with a virtual barrier. Chapter 7 applies describing function analysis from nonlinear controls theory to predict limit cycles for a haptic knob contacting a virtual barrier. Chapter 8 presents empirical observations of limit cycles in a haptic knob contacting a virtual barrier, for various displacement quantizations and sampling rates, which agree with the simulation results. Chapter 9 discusses the effects of sample rate and displacement quantization on

the amplitude and frequency of limit cycle oscillations. Chapter 10 investigates the implications of this work for the design of haptic devices. Chapter 11 discusses potential approaches to mitigating limit cycle oscillations. Chapter 12 concludes the dissertation. Several appendices follow with supporting material, including calibration descriptions and some raw data.

Chapter 2 : Human System Identification Methods and Literature

Simulation and theoretical analysis of control systems requires a model of the physical plant, in this case the haptic knob and the human hand grasping it. We have two choices for obtaining the plant model: analytic construction and system identification. Analytic construction (building a theoretical physical model from precise knowledge of a system's characteristics) is best-suited to simple systems about which the investigator has accurate knowledge. System identification is an iterative process of experimentation and analysis, where the system is presented with an exciting input, the output is measured, and analysis of the relationship between the input and output is used to generate estimates of the plant model. System identification is useful for more complex systems, especially if some simplifying assumptions can be made. The human hand, with at least 21 degrees of freedom, is just such a complex system. A hand grasping a haptic knob with very small oscillatory displacements offers an example of a complex system with relatively constrained motion that is well-suited to system identification.

The area of human systems identification that has received the most careful attention is the identification of human joint dynamics. Common system identification techniques for the study of human joint dynamics include both the study of commanded force/torque

trajectories and their resultant position/angle vectors, as well as the study of commanded position/angle trajectories and their resultant force/torque trajectories. Ideally, the two approaches should yield comparable results, but Kearney and Hunter assert that estimates of joint dynamics obtained with torque inputs will be better at low frequencies than high, and estimates obtained with position input will be better at high frequencies than low (Kearney and Hunter, 1990). To help compensate for lower response at high frequencies using torque input, they suggest using a torque input with more power at high frequencies. They caution that the position-input approach is technically more difficult, since “the actuators used must have a higher dynamic response and larger force-generating capacity than the joint under study in order to achieve effective position control.”

The obvious approach for a haptics engineer familiar with torque motors and optical encoder rotation sensors would be to use a torque-input approach, equating torque with the commanded current input to the motor, and measuring angular position with an optical encoder. This arrangement avoids the expense, complexity, and inertia of a torque sensor; however, Kearney and Hunter caution against this approach, stating that the current-torque relationships of motors have dynamics, nonlinearities, and time-varying characteristics that can lead to serious system identification errors. They assert that independent measurement of torque is essential (Kearney and Hunter, 1990).

Previous studies of human dynamics have reported wide intersubject variability. This corresponds to the haptics engineer’s intuition that devices that exhibit limit cycles with one individual may not exhibit limit cycles with another individual, or may have a limit cycle but at a different frequency. Kearney studied human ankle dynamics,

reporting that inertia, elasticity, and the slope of the linear relation between torque and joint stiffness, K , had an intersubject coefficient of variation between 30% and 50% (Kearney et al., as cited in Kearney and Hunter 1990, ref. 137). Interestingly, there was substantially less variability in the damping parameter, with an intersubject coefficient of variation of only 15%. Kearney and Hunter used a test-retest paradigm for their study, and reported excellent intrasubject repeatability. This is a heartening result for experimentalists familiar with the difficulties of recording data from humans, but low intrasubject variability requires diligence from the experimenter in avoiding the many threats to repeatability.

Many sources have supported the validity of modeling human joint dynamics as linear about an operating point, with second-order models being common (Agarwal and Gottlieb, 1977; Becker and Mote, 1990; Crowninshield et al., 1976; Gillespie et al., 1999; Hogan, 1990; Hunter and Kearney, 1982; Kearney and Hunter, 1982, 1983, 1990; Kearney et al., 1997; Milner and Franklin, 1995, 1998). Most modeling approaches assume a time-invariant system – expanding the scope to include time-varying behavior greatly increases model complexity and the repeatability challenge. For a human hand in a haptic knob grasp, this implies that the grip remains constant, motion about the origin is small to avoid changing kinematics, and muscle activation does not change to affect the stiffness or damping characteristics of the grasp.

The potential for changes in muscle activation raises some challenges to time-invariance. Muscle stretch reflex responses can be seen in EMG signals from the hand muscles in as little as 20-30 ms (Milner and Franklin, 1995). By comparison, muscle stretch reflexes for the ankle have been shown to have a minimum latency of 40 ms

(Kearney et al., 1997). Cutaneous slip reflexes can occur in fingers grasping an object at about 70 ms after onset of slip (Johansson and Westling, 1984). Voluntary muscle activation occurs at longer latencies. One strategy to avoid the complication of changes in muscle activation is to apply and remove input stimuli rapidly, before any voluntary or reflexive muscle activation can occur (Hajian, 1997; Kearney and Hunter, 1990; others). This is a valuable approach, but limits the number of input waveforms that one can use. Another approach would be to simply instruct the subjects not to change muscle activation. For example, in the case of a haptic knob, subjects could be instructed to apply a “light grasp” or a “heavy grasp.” This approach might have some limited value, but risks high intrasubject variability.

A grip force sensor could be added to the knob, and the subject could be instructed to maintain a particular grip force during the application of the input stimuli (receiving visual feedback of target and actual grip forces on a computer monitor). This would allow some level of control and measurement of muscle activation, and would allow for the examination of the relationship between grip force and grasp impedance; however, it only achieves partial control of the variability. By contracting agonist and antagonist muscles at the same time, individuals can vary the impedance of their joints without changing the resultant torques (forces) (Hogan, 1990; Kearney and Hunter, 1990). The level of independent control individuals have over cocontraction likely varies from joint to joint, and may not be as high when the subjects are exerting high forces (Kearney and Hunter, 1990). Levels of cocontraction might be observed with electromyographical (EMG) measurements of muscle activity, at the cost of considerable experimental complexity.

2.1 System Identification Input Waveforms

System identification involves the application of an input stimulus to a system, and the measurement of the system's response. Appropriate selection of an input stimulus plays an important role in the quality of the resulting system identification. Many types of input stimuli exist. Factors such as system dynamics and noise levels influence the relative performance with different input stimuli. System identification of human limb dynamics adds another dimension to the problem, since reflexes and voluntary muscle activation can alter system dynamics during presentation of all but the shortest input stimuli. Kearney and Hunter provide an excellent review of the characteristics of various input waveforms for the purposes of human limb dynamics identification (Kearney and Hunter, 1990). Input stimuli fall into three categories: transient, sinusoidal, and stochastic.

Transient – This category includes infinite impulses, finite pulses (e.g., Hajian, 1997), steps (e.g., Milner and Franklin, 1995, 1998), and ramps. Theoretically, infinite impulses provide a perfect stimulus with even representation of all possible frequencies, and can be used to determine nonparametric models in noise-free systems. In practice, only finite pulses can be attained. Finite pulses tend to fare poorly in the presence of noise, and offer much less power at higher frequencies. While transient inputs are least appealing for nonparametric models, many investigators have attempted to determine parametric models of joint dynamics using transient inputs. Hajian's work using 20 ms pulses to determine finger dynamics is particularly relevant (Hajian, 1997). Transient inputs, if short enough in duration, have the benefit of avoiding dynamic changes due to reflexes or

voluntary responses. For parametric system identification of human joint dynamics, transient inputs such as pulses are an attractive option.

Sinusoids – Sinusoids offer an intuitively appealing method of system identification.

Using a series of individual sinusoids, one can determine the gain and phase response of a system across a range of frequencies. This approach concentrates all the input power at a given frequency, maximizing the signal-to-noise ratio. Measuring response power at frequencies other than the driving frequency can help to reveal nonlinearities. Selection of particular frequencies can help to explore regions of interesting system behavior in detail. Finally, analysis using this approach can be simple; however, sinusoids have four significant disadvantages. First, data acquisition is time-intensive and results in a large amount of data that one must store and analyze. Using a “chirp” input can improve this situation, though the chirp must be applied with caution to avoid transient effects.

Second, the results are only valid if the system does not change during the duration of the experiment with all the individual sinusoidal trials – this is next to impossible with human subjects. Third, using lengthy stimuli such as sinusoids makes it difficult to eliminate the effects of reflexes and voluntary muscle activation. Fourth, sinusoidal inputs have been reported to evoke some strange, nonlinear responses in human joint studies (Kearney and Hunter, 1990).

Stochastic – These inputs include random noise and pseudorandom signals. The most popular stochastic inputs are Gaussian white noise (e.g., Agarwal and Gottlieb, 1977; Becker and Mote, 1990) and pseudorandom binary signals (e.g., Hunter and Kearney, 1982, 1983; Kearney and Hunter, 1982). Both contain power over a wide frequency range, though pseudorandom binary signals have more power for a given input

amplitude. Gaussian noise inputs have another disadvantage: since they can never be perfectly white, more input trials will be required, leading to a larger data set.

Pseudorandom signals result in relatively small data sets, and have benefits that make them particularly suited to system identification of human joint dynamics – since they are unpredictable, subjects cannot voluntarily respond to them in a consistent way, and responses from several presentations of the same input can be averaged to reduce the effects of noise and random voluntary responses. For these reasons, pseudorandom binary sequences are quite attractive for system identification of human joint dynamics.

2.2 Previous Studies of Finger Dynamics

Several studies have been conducted on the dynamic properties of the human finger. Hajian conducted the study most relevant to the current work. A thorough description of Hajian's work appears below. Becker and Mote (1990) studied the dynamics of the index finger in abduction/adduction, finding that a second order mass-spring-damper model described the dynamics well for small displacements, and that fatigue of the finger muscles reduced the stiffness and damping parameters. Milner has conducted studies on fingers specifically motivated by haptic feedback applications (Milner and Franklin, 1995, 1998). Karason and Srinivasan (1995) studied finger dynamics in a grasp of an active instrumented object that could rapidly contract or expand, deriving a third-order model with separate terms for the finger impedance and fingerpad impedance. Gulati studied the *in vivo* compressibility of the human fingerpad (Gulati, 1995; Gulati and Srinivasan, 1995; Srinivasan et al., 1992). Further work in the same laboratory has examined the viscoelastic properties of the fingerpad (Birch and Srinivasan, 1999).

Pawluk and Howe have also investigated the dynamics of the fingerpad in compression (Pawluk and Howe, 1999a,b).

For his doctoral dissertation in Howe's Harvard BioRobotics Laboratory, Hajian conducted an extensive study of the impedance of the human fingers (Hajian, 1997). He began with an investigation of the impedance of the straightened index finger in extension and abduction at the metacarpal-phalangeal (MCP) joint, continued with an investigation of the impedance of a pinch grasp, built on these results to present and validate a model of the human hand in a drum roll on a musical drum, and finished with an implementation of robotic drumming. Hajian showed that drummers are able to overcome bandwidth limitations on active control by modulating their passive impedance, creating drum rolls at a higher frequency than they could control actively.

In his first study examining the impedance of the index finger MCP joint in extension and abduction, Hajian used transient forces with a maximum duration of 20 milliseconds to avoid the onset of the stretch reflex at approximately 30 milliseconds, and to avoid cutaneous slip reflex (Johansson and Westling, 1984) and voluntary muscle contraction. He used five subjects holding a rigid handle and extending their finger over a pneumatic piston with a force sensor and piezoelectric accelerometer. The subjects began slowly pressing against the piston, and at a force level set by the experimenter (between 2 N and 20 N), the piston would apply a transient force displacing the finger approximately 5 mm.

Hajian assumed a linear, second-order translational model at the fingertip, with parameters m , b , and k :

$$m\mathbf{x}''(t) + b\mathbf{x}'(t) + k\mathbf{x}(t) = \mathbf{f}(t)$$

Hajian measured applied force and finger tip acceleration just prior to expansion of the cylinder, and calculated velocity and displacement curves from the acceleration signal. Writing the above equation in matrix form allowed Hajian to divide the position, velocity, and acceleration matrix by the force vector to obtain the m , b , k parameter vector (using the matrix division function in MATLAB gives a least-squares fit). Results showed increasing damping and stiffness parameters for all subjects with increasing finger bias force – as subjects pressed harder against the plate, their fingers exhibited more damping and stiffness. Hajian also computed the damping coefficient, finding that subjects' fingers were critically damped at bias force levels over 4N (supporting anecdotal observations by many that light grips are more prone to instability). Mass estimates remained relatively constant, except at low bias force levels. Hajian attributed this variation to increased interphalangeal (distal and middle knuckle) joint compliance in the relatively relaxed fingers that led to non-rigid-body motion of the finger. Hajian notes a linear relationship between bias force and stiffness, but observes that stretch reflexes and co-contraction will probably act to increase stiffness, warning that his estimates obtained without these effects likely represent a lower bound on stiffness.

Hajian's second study measured variations in finger impedance in a pinch grasp of a freely supported rigid object. Subjects grasped the object's handle between the palm, middle, ring, and little fingers while pinching another part of the object between the thumb and forefinger. Hajian performed the tests using three male subjects, and pinch grasp forces of 8, 20, and 30 N, with subjects regulating pinch grasp force by visually matching a target force on an oscilloscope. Subjects performed the test at palm grasp force levels of 20, 50, and 80 N for each of the pinch grasp force levels.

Hajian computed a separate second-order lumped parameter model for each finger in a manner similar to that of the first experiment. Results were similar, though he found only 87% variability accounted for (VAF) in the second experiment as compared to 97% for the first experiment. Mass estimates for subjects' thumbs and forefingers did not vary significantly over the range of finger grasp forces and palm grasp forces. Both damping and stiffness increased almost linearly with finger grasp force. Hajian found that damping and stiffness estimates increased slightly with increasing palm grip force, implying that cocontraction of the muscles in or near the thumb and forefinger was contributing to stiffness and damping.

Chapter 3 : Apparatus


Investigating the dynamics of the human hand grasping a haptic knob requires a testbed with a high-fidelity torque actuator having as little friction and inertia as possible. A position sensor should have high enough resolution to give excellent velocity signals and to far exceed the resolution available on commercial haptic knobs (currently in the range of 4,000 optical encoder counts per revolution, after quadrature). Direct measurement of rotational acceleration would be beneficial, but was not implemented for the current work.

3.1 Hardware Setup

To meet these requirements, a haptic knob reference system dubbed “HREF” was designed and built (see Acknowledgements for design team credits). Figure 1 shows a 3D CAD drawing of the system. HREF uses a Maxon RE025-118752 motor, similar to those used on haptic feedback interfaces from Immersion Corporation and SensAble Technologies. A TA-115 transconductance amplifier from Trust Automation, San Luis Obispo, CA, drives the RE025 motor. One D/A channel of a National Instruments 6052E PCI I/O card drives the command input of the TA115 motor amplifier.

HREF possesses a 640,000 count per revolution reflective-diffraction “Mercury” optical encoder from MicroE Systems Corporation (MicroE, 2001). The HREF design has provisions for a disk mounted on the motor shaft to carry a small accelerometer near the outer edge of the disk with the measurement axis aligned tangent to the disk (using a linear accelerometer to measure angular acceleration), though this feature has not been implemented. Other measurement capabilities include knob torque sensing and the “squeeze” force of the human subject’s grip on the knob. The latter has been implemented for several different-sized knobs, using an Entran EFLS-B1 100 N (22.5 lb) button-style compression load cell amplified by an Entran PS30A amplifier with outputs connected to differential A/D inputs on a National Instruments 6052E I/O card.

Figure 1 contains an illustration of the HREF apparatus. The illustration includes a reaction torque load cell in the center foreground that was tested with the system, though off-axis moments on the load cell prevented reliable torque measurements during the present experiments.



Design and drawing: B. Schena

Figure 1: Haptic Reference (HREF) knob apparatus

3.2 Software Architecture

Two computers were used during data collection for all experiments. A host computer ran a Matlab shell program under a Windows operating system, and a target computer ran a real-time executable program over a real-time operating system kernel provided as part of the Mathworks Real-Time Workshop xPC Target system. The target executable began before run-time as a Simulink model program on the host, and was compiled for the target and downloaded to the target over a TCP/IP connection. This architecture created some interesting capabilities, as when investigators downloaded a target program from a host in the San Francisco Bay Area to a target PC in Montreal to control a real-time haptics application. This cross-continent collaboration proved useful for debugging similar laboratory setups.

The target PC contained the National Instruments 6052E interface card with differential A/D lines for reading the torque and grip load cells, and D/A for outputting command signals to the motor amplifier. The target PC also contained a CIO-QUAD04 board from Measurement Computing, Inc. (formerly ComputerBoards) for reading quadrature encoder signals from the optical encoder rotation sensor on HREF. The target program ran for the duration of a block of trials, reading sensors and commanding motor torque at a rate of 10 kHz, with periodic parameter updates over TCP/IP from the host computer running the Matlab shell program to cycle through experimental conditions (grip force threshold, pulse strength, etc.). A state machine running on the target system sensed grip threshold and determined when to fire a pulse, with parameters set by the host. During conduct of a block of experimental trials, data was automatically uploaded to the host after each pulse to be stored for later analysis. The shell program on the host

also displayed the data to the experimenter as each pulse occurred to ensure that data collection was proceeding without complication.

Chapter 4 : System Identification for a Fingertip Pinch Grasp

4.1 Introduction

Dynamics for a hand grasping a knob will vary considerably from user to user, for various grasp postures, with different grip strengths, and potentially with other variables. This chapter describes system identification of the human hand in a fingertip pinch grasp, with the thumb and forefinger in opposition, pinching the knob, as in Figure 2.

4.2 Methods

4.2.1 APPARATUS AND PROCEDURE

This experiment used the apparatus described in Chapter 2. Nine healthy subjects between the ages of 23 and 32 participated; five male and four female. For each trial, an audible “beep” signaled subjects to start slowly squeezing the knob in a pinch grasp. When the subject’s grip strength reached a preset threshold for a given trial, the HREF knob applied a 20 ms clockwise torque pulse to the subject’s hand, displacing it approximately 0.07 – 0.11 radians. The subject then relaxed his or her hand to a preset level and awaited the next “beep” from the computer. Pulse magnitudes were increased

with increasing grip strengths so that the rotational displacement of the knob remained roughly the same across trials for different grip strengths.

Subjects completed the tests in blocks of 18 trials, with three trials in each block for each of six grip strengths: 0.72, 1.3, 2.4, 4.3, 7.8, and 14.2 Newtons. The trials in each block were presented sequentially, starting with three trials at the lightest grip strength and progressing to three trials at the strongest grip strength. Subjects completed one training block of trials, and three experimental blocks for which data was analyzed. Prior to beginning the training block, subjects were instructed to assume a neutral and relaxed posture, with chair height adjusted so that the right forearm remained close to level, and the fingers of the right hand grasping the knob in the pinch grasp illustrated in Figure 2. Subjects were asked not to allow their hand or arm to touch the laboratory table top upon which the knob rested, and to avoid significant extension or flexion of their wrists.

Subjects grasped the knob with fingerpads centered on red tape strips indicating the center of the grip force load cell, and kept their fingers as close as possible to directly opposite each other along an axis passing through the center of the load cell. The desired grasp posture was designed to be realistic, repeatable, and to keep the grasp axis passing through the grip load cell for accurate grip force measurement (see Appendix B for a description of grip load cell calibration and its limitations). Subjects were coached on grasp posture during the training block of trials, and subjects did not have trouble maintaining a grasp that appeared satisfactory to visual inspection throughout the blocks of trials. Subjects were instructed to keep their grasp as constant as possible through a given block of trials, but were encouraged to let go of the knob and move freely to relax between blocks.




Figure 2: Subject's hand grasping HREF knob in a pinch grasp.

4.2.2 DETERMINATION OF THE RELEVANT DISPLACEMENT RANGE

If one wants to obtain dynamic model information for the human hand in a knob grasp in order to use that information for simulation and prediction of instability, it makes sense to determine the operating points at which instability is likely to occur. What angular displacements and torques are associated with the unwanted "buzzing" of a haptic knob? Since human limbs possess nonlinear dynamics and can only be considered linear about a given operating point, the determination of relevant operating points is important. A simple test was performed with one human subject to obtain an approximation for the magnitude of limit cycle oscillations. Figure 3 shows a plot of the angular displacement of a haptic knob as the subject turns the knob between three detents (composed of piecewise continuous haptic springs). The haptic knob had 12 detents per revolution, and one can observe about 0.52 radians (6.28 radians/12) between detents levels in Figure 3. The oscillations have a magnitude of approximately 0.16 radians peak-to-peak (+/- 0.08

radians), a frequency of 33 Hz, and are associated with oscillatory torque commands of approximately +/- 6 mNm. Peak detent forces were approximately 8 mNm (8×10^{-3} N-m) for this haptic knob test run. The subject anecdotally experienced these oscillations as rather strong but not inconsistent with experience. "Haptic buzzing" of smaller amplitudes may be of interest as well.

Note the greatly diminished oscillations during the second detent transition in Figure 3 (just after Time = 1 second). During the trial, the subject was able to completely prevent oscillations, or eliminate them once they started, by squeezing the knob harder, presumably increasing the damping and stiffness coefficients of his grasp.

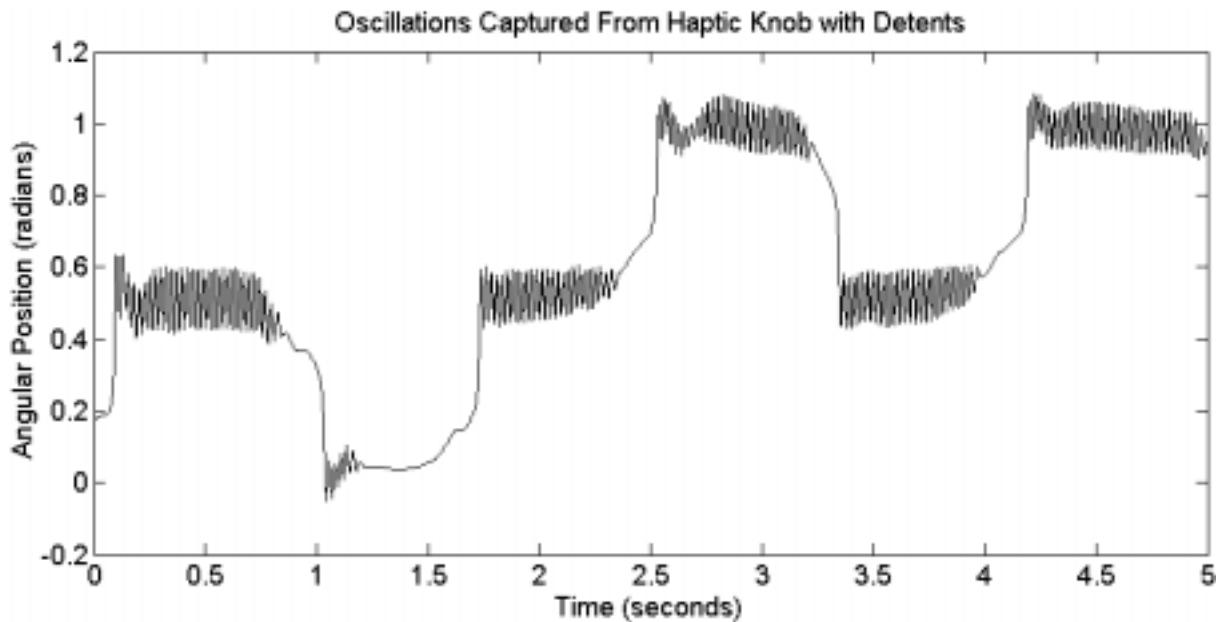


Figure 3: Oscillations Captured from Haptic Knob with Detents

4.2.3 SYSTEM IDENTIFICATION TECHNIQUE

This experiment uses similar assumptions and technique to that of Hajian (Hajian, 1997). It assumes a linear, second-order rotational model at the fingertip, with

parameters J , B , and K representing the moment of inertia, damping coefficient, and stiffness coefficient, respectively:

$$J\theta''(t) + B\theta'(t) + K\theta(t) = \mathbf{T}(t)$$

Figure 4 contains a schematic representation of the lumped second-order model:




Figure 4: Second-order lumped parameter model

The experimental apparatus collected raw data for knob position, commanded torque, reactive torque measurement, and grip force measurement. Estimates of velocity and acceleration were obtained by differentiation, and double differentiation, respectively, of the rotational position signal. Prior to differentiation, raw displacement data was averaged for all trials for the same condition in a block. The position and velocity signals were smoothed by filtered decimation and interpolation. The high sample rate (10,000 samples/second), extremely high encoder resolution (640,000 counts/revolution) and relatively noiseless encoder signal made double differentiation a viable approach. Writing the above equation in matrix form allows division of the rotational position, velocity, and acceleration matrix by the torque vector to obtain the J , B , K parameter vector (using the matrix division function in MATLAB gives a least-squares fit). Since data was captured for both commanded torque and measured torque, either data vector could be used in the estimation of J , B , and K .

4.3 Results

Figure 5 contains thumbnail summary plots of displacement trajectories for all 54 trials for one subject. The trajectories have a significant variability in peak magnitude, but all trajectories for a given grip strength have consistent shapes. Figure 6 contains a composite plot showing commanded torque (upper left), acceleration (upper right), velocity (lower left), and position (lower right) vectors. The onset of the pulse causes a sharp acceleration pulse, with steadily ramping velocity and increasing position displacement. Figure 7 contains a plot of commanded torque, estimated torque contributions from moment of inertia, damping, and displacement, as well as total estimated torque. The estimated inertial, damping, and stiffness torques represent the product terms on the left-hand side of the equation:

$$J\theta''(t) + B\theta'(t) + K\theta(t) = \mathbf{T}(t)$$

obtained by multiplying estimated J , B , and K parameters by the acceleration, velocity, and displacement trajectories, respectively. For a perfect model estimate, the estimated inertial, damping, and stiffness torque contributions would add to create a total estimated torque vector that equals the commanded torque.

Figure 8 through Figure 11 show the results of the experiment. Figure 8 shows measured step responses for six different grip forces in a given trial (solid lines), simulated trajectories based on model estimates derived from the same data (dotted lines), and simulated trajectories based on models derived using data from a different trial (dot-dashed lines). This example shows excellent agreement between the data and a model based on the data, and good agreement between the data and a second model estimate based on different data from the same subject under the same conditions.

Figure 9 shows moment of inertia, damping, stiffness, and damping ratio estimates for one subject across all three blocks of trials, for all six grip strengths. The equation $\zeta = \frac{B}{2\sqrt{JK}}$ defines the damping ratio for the system. Error bars indicate standard deviation across the three blocks of trials. Figure 10 presents moment of inertia, damping, stiffness, and damping ratio estimates for all nine subjects. These plots for all subjects reveal quite similar trend behavior for each parameter. Figure 11 contains plots of moment of inertia, damping, stiffness, and damping ratio estimates averaged across all subjects, with error bars indicating standard deviations. All plots in the three previous figures include the lumped parameter estimates for the combined hand-apparatus system. Apparatus stiffness, damping, and moment of inertia were empirically determined using similar methods to those used for human subjects. Stiffness and damping inherent in the HREF apparatus were negligible, and a moment of inertia of 19.8 gm-cm² was determined (see Appendix C for apparatus dynamics identification details).

Between-subjects coefficients of variation were calculated using the data presented in Figure 11, and the coefficient of variation equation:

$$CV = \frac{\sigma}{\mu},$$

where σ represents the standard deviation of the parameter estimates across subjects for a particular grip force, and μ represents the mean parameter value across subjects for a particular grip force. Coefficients of variation for moment of inertia ranged from 4% to 10% (though since moment of inertia is dominated by the apparatus, this represents a much larger percentage of the finger inertia). Coefficient of variation for damping ranged

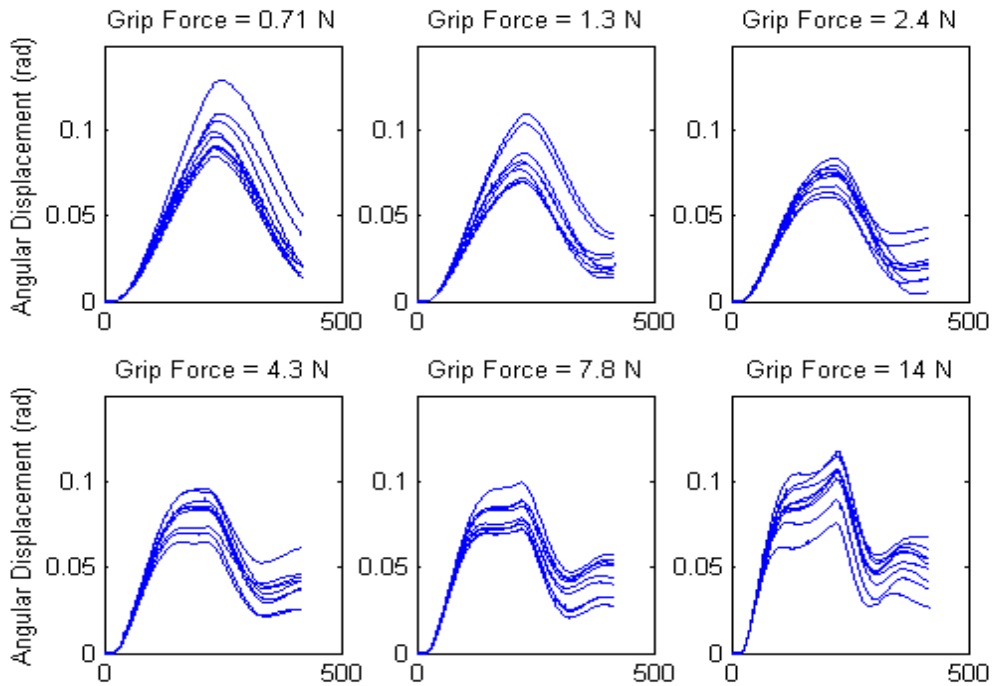


Figure 5: Thumbnail summary plots of displacement trajectories for all 54 trials for a particular subject; 9 trials per grip force condition, with increasing grip force from left-right and top-bottom.




Figure 6: Commanded torque (upper left), acceleration (upper right), velocity (lower left), and displacement vectors (lower right) for a typical trial with a grip force of 2.4 N.




Figure 7: Commanded torque (T_{q_cmd}), total estimated torque (T_{q_calc}), and estimated torque contributions from moment of inertia (T_{q_acc}), damping (T_{q_vel}), and displacement (T_{q_disp}) [plot layout adapted from Hajian, 1997].

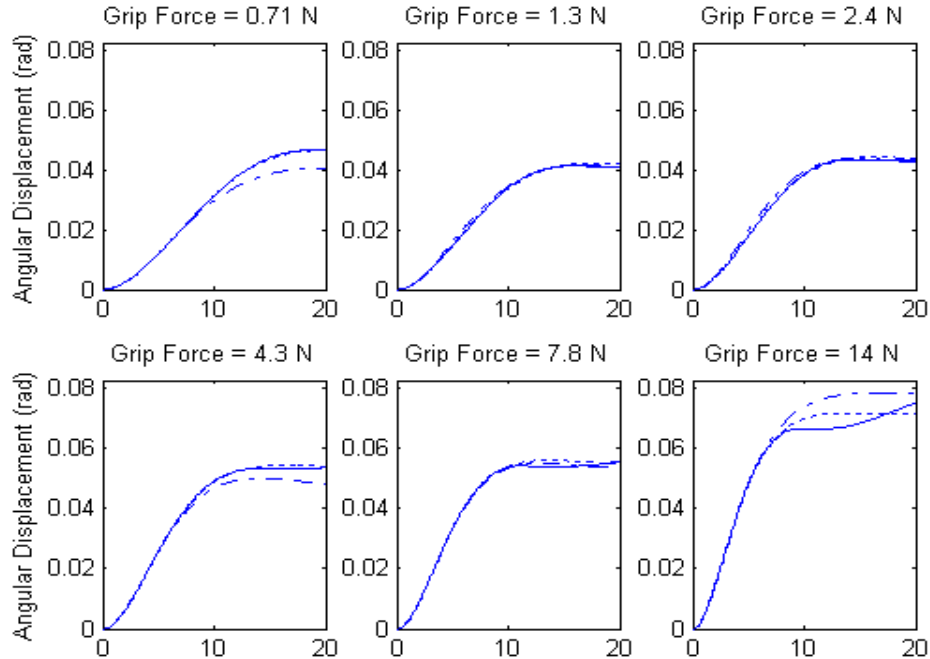


Figure 8: Measured step responses for a trial at six different grip forces (solid lines), simulated trajectories based on model estimates derived from the same data (dotted lines), and simulated trajectories based on models derived using data from a different trial (dot-dashed lines).




Figure 9: Moment of Inertia, Damping, Stiffness, and Damping Ratio for One Subject

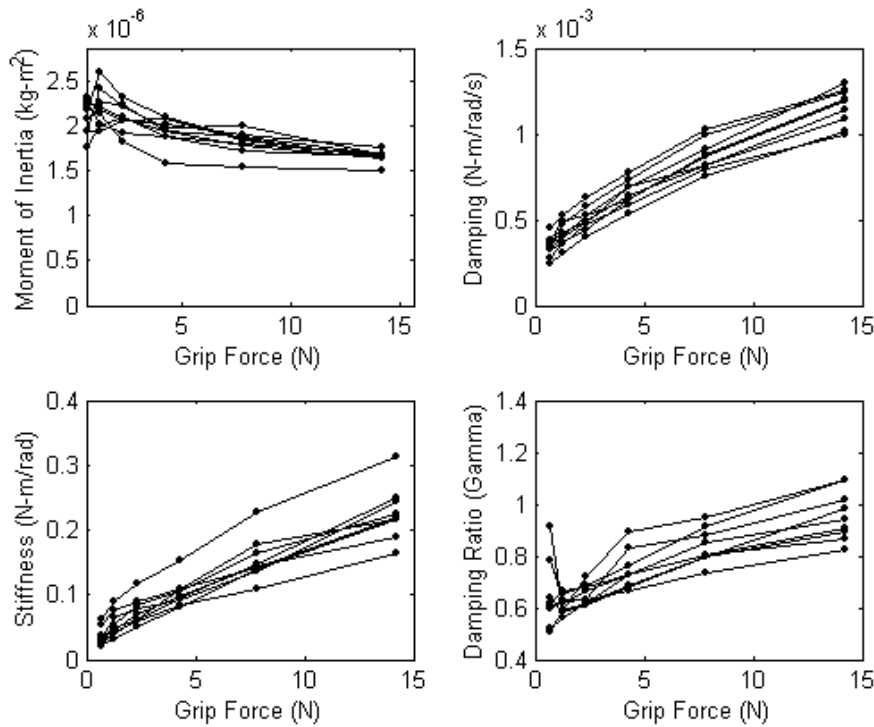


Figure 10: Moment of Inertia, Damping, Stiffness, and Damping Ratio for All Subjects

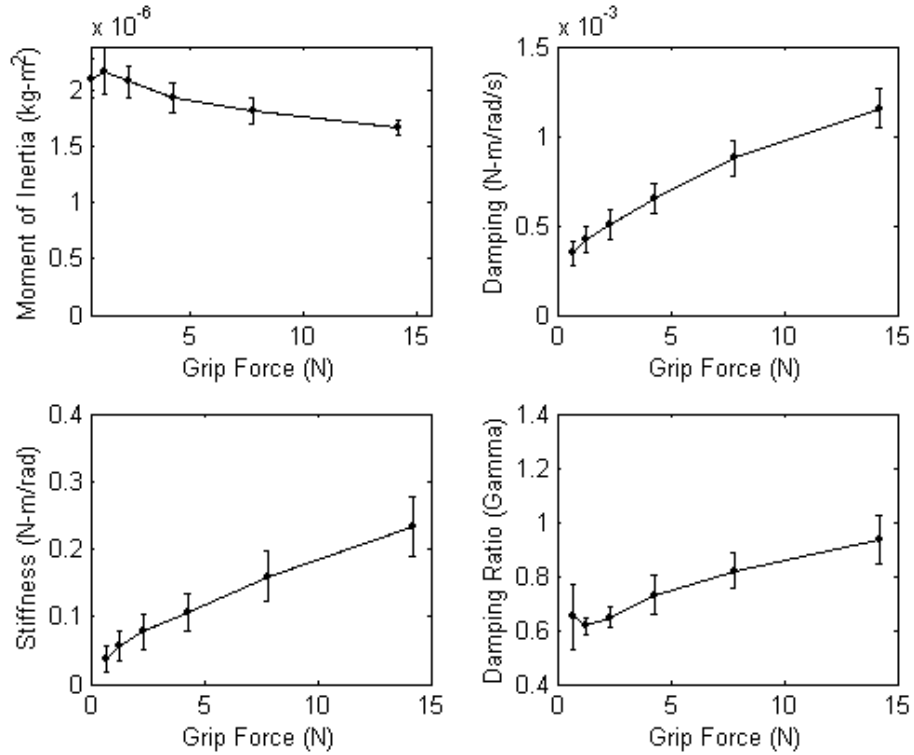


Figure 11: Moment of Inertia, Damping, Stiffness, and Damping Ratio Averaged Across All Subjects

Table 1: Results for Experiment 1, Simple Pinch Grasp on 17.8 mm Knob, Across Subjects

Grip Force (N)	J (kg-m ²)			B (N-m/rad/s)			K (N-m/rad)			ζ zeta		
	J	std J	CV J	B	std B	CV B	K	std K	CV K	zeta	std zeta	CV zeta
0.71	2.11E-06	1.71E-07	8.1%	3.45E-04	6.47E-05	18.8%	0.04	0.02	48.3%	0.6	0.12	18.6%
1.3	2.18E-06	2.05E-07	9.4%	4.20E-04	7.55E-05	18.0%	0.06	0.02	39.0%	0.6	0.03	4.9%
2.4	2.08E-06	1.41E-07	6.8%	5.09E-04	8.33E-05	16.3%	0.08	0.03	33.8%	0.6	0.04	6.0%
4.3	1.94E-06	1.35E-07	7.0%	6.53E-04	8.25E-05	12.6%	0.11	0.03	25.4%	0.7	0.07	9.8%
7.8	1.82E-06	1.16E-07	6.3%	8.76E-04	9.75E-05	11.1%	0.16	0.04	23.3%	0.8	0.06	7.9%
14.2	1.67E-06	6.83E-08	4.1%	1.15E-03	1.08E-04	9.3%	0.23	0.04	19.1%	0.9	0.09	9.7%

from 9% to 19%. Coefficient of variation for stiffness ranged from 19% to 48%.

Coefficient of variation for damping ratio ranged from 5% to 19%.

4.4 Discussion

Results for damping and stiffness match the experience of haptic researchers and designers – increasing grip strength increases stiffness and damping. The nearly linear increase in damping and stiffness with increasing grip force compares well to results of other researchers (e.g., Hunter and Kearney, 1982, 1983; Kearney and Hunter, 1982; Hajian, 1997). This behavior allows one to eliminate some unwanted vibrations in a haptic system by gripping the manipulandum with a stronger grip, and makes a light grasp the preferred grasp of persons attempting to excite “buzzing” limit cycles in a haptic system. In fact, a common approach to excite limit cycles with some mechanisms (such as a two-dimensional pantograph) is to tilt them slightly so that the manipulandum moves with a small force against a virtual wall, with absolutely no additional damping or stiffness from a human grasp. The parameter estimates presented in this work may be useful to haptic designers seeking to simulate new system configurations with the inclusion of a “virtual human” in the simulation.

Assuming a constant grasp posture and fixed coupling between the fingers and knob, moment of inertia should remain constant for each subject across all grip strengths, though it appears to decrease with increasing grip strengths in the present experiment. This apparent and unexpected sensitivity of moment of inertia to grip force is not a great concern for several reasons. First, the change in moment of inertia with increasing grip force is small compared to the marked changes in stiffness and damping. Second, the

present experiment spanned a wide range of grip forces. The moment of inertia estimates seem most suspect at high grip forces; however, the important system identification results occur at light grip forces which are associated with the worst limit cycle oscillations. Third, subsequent simulations and nonlinear system analysis using the models developed here supports their value for their intended purpose. Despite the relative unimportance of the variability in moment of inertia measurements, potential reasons for this variability were investigated in an effort to provide a more complete understanding of the results.

Hajian found that changes in finger posture with increasing force led to differences in estimated inertia for the finger in extension (Hajian, 1997); however, differences in grasp posture were not observed as grip force increased. Consideration in the present work of a fourth order system consisting of two cascaded second order systems for the finger and for the fingerpad/apparatus provides a possible explanation for the decreasing moment of inertia estimates observed with increasing grip force (see Page 39). Karason and Srinivasan (1995) confronted a similar situation and developed a third-order model that also treats the finger and fingerpad separately. Another possible contributing factor is the decrease in effective radius as the fingerpad becomes more compressed with higher grip forces and the mass of the finger and thumb move closer to the center of rotation around the knob axis.

One additional consideration that may have had an impact on moment of inertia estimation regardless of grip force is the concept of implicit regularization of least-squares estimates. If a least-squares fitting routine attempts to match several trajectories (acceleration, velocity, and displacement, for example) and one of the trajectories is

rather error prone (e.g., acceleration) while the others are relatively reliable (e.g., velocity and displacement), the least-squares fitting routine will tend to minimize the coefficient (e.g., moment of inertia) for the least reliable trajectory in an effort to minimize error.

The equation below represents the extreme of this situation, where an error term is added to acceleration, and the velocity and displacement trajectories are perfect. In an effort to minimize the error in the estimate of $\mathbf{T}(t)$, the least-squares algorithm will tend to underestimate J .

$$J[\boldsymbol{\theta}''(t) + \boldsymbol{\varepsilon}(t)] + B\boldsymbol{\theta}'(t) + K\boldsymbol{\theta}(t) = \mathbf{T}(t)$$

Damping ratio estimates for the current experiment reveal underdamped behavior in almost every case, with the damping ratio approaching 1.0 only for the strongest grip forces. Knob rotation for the simple pinch grasp studied here primarily involves finger abduction/adduction. Two studies of finger abduction/adduction dynamics show underdamped behavior to be more likely for this case (Becker and Mote, 1990; Hajian, 1997). These observations for abduction/adduction are consistent with the results of the current experiment. The underdamped behavior for abduction/adduction and knob rotation with a pinch grasp differ from findings for other situations such as finger flexion/extension, where the behavior is more likely to be overdamped or close to critically damped (Hajian, 1997).

Intrasubject variability for all four parameters was moderate, as was intersubject variability. (Error bars in Figure 9 indicate intrasubject variability for one selected subject, and the error bars in Figure 11 indicate intersubject variability.) Data from all subjects showed similar trends for each of the four parameters with changes in grip force. Data from all subjects indicate a significantly underdamped system for all but the

strongest grip force. Subjects commented that this grip force was much stronger than they would typically use to grasp a knob (they also commented that the lowest grip force was much lighter than they would use in practice). The underdamped system estimates contrast with the results of Hajian (1997), whose subjects were more likely to exhibit close to critically damped behavior for finger extension experiments.

Figure 12 contains a comparison of normalized damping and stiffness results for the present experiment to results obtained by Hajian for his finger extension and abduction experiments. Grip/finger force vectors (x-axis), and damping and stiffness coefficient vectors, were each normalized by dividing by the midpoint (average of the minimum and maximum values) of the vector. This normalization allows one to inspect the relative slopes and curve shapes despite the fact that independent and dependent variables from the three experiments do not share the same ranges or physical units (Hajian used Cartesian units, the present experiment uses angular units). Error bars represent ± 1 standard deviation, and have been similarly normalized. Caution should be used in interpreting Hajian's results as graphed here; he only presented his results in tabular form for two or three grip forces, so the curves plotted here represent two- (for abduction) and three- (for extension) point lines. All three experiments (the present one and two from Hajian) show increasing damping (left plot in Figure 12) with increasing grip/finger force, and all three show quite similar increases in stiffness with increasing grip/finger force (right plot in Figure 12). The present experiment, with thumb and forefinger in a pinch grasp of a knob perturbed by a torque pulse, involves the thumb and forefinger in abduction/adduction, with the fingerpads in shear. Hajian's abduction experiment studied

the finger in abduction, but involved compressive contact with the side of the finger's distal phalanx.

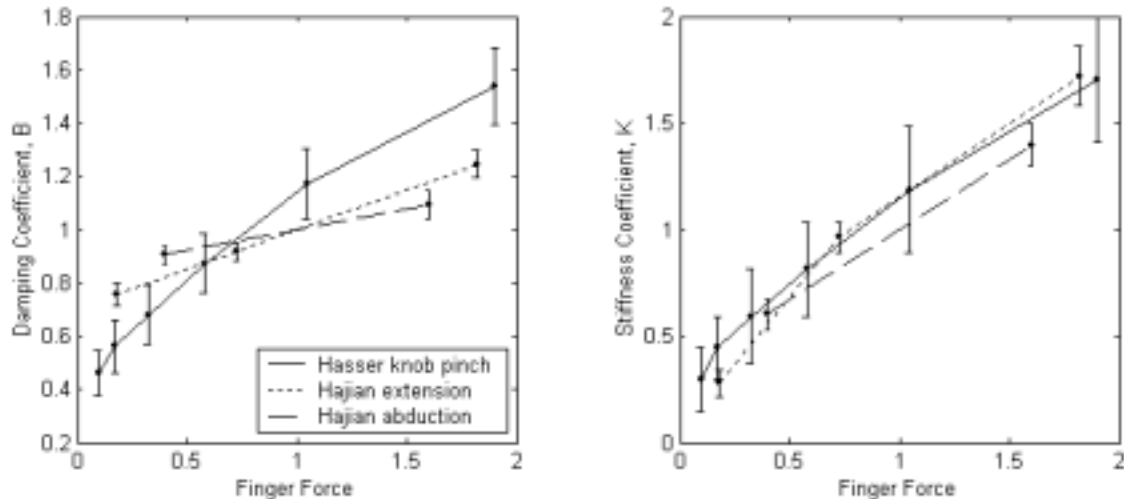


Figure 12: Comparison of Normalized Results to Those of Hajian's Finger Extension and Abduction Experiments

The results presented here are valid for the specific displacement magnitudes that occurred during this experiment. Isolated muscle typically exhibits higher stiffness when subjected to small displacements than when subjected to large displacements, and one study found significant decreases in ankle stiffness for larger displacements (Kearney and Hunter, 1982).

The sequential ordering of grip force conditions during data collection for this experiment, from lightest grip to strongest grip in each block, raises a question about the potential for muscle fatigue to introduce systematic artifacts into the data. This concern is mitigated by the fact that a block of trials typically lasted only 3-4 minutes, and was accompanied by a 1-2 minute rest before the next trial. A previous study found that fatiguing of the ankle for as long as 80 seconds did not alter ankle dynamics (Hunter and Kearney, 1983). A study of fatigue effects on finger abduction/adduction dynamics (quite relevant to the current experiment) showed significant fatigue effects; in the

fatigued case, finger stiffness and damping decreased markedly, reducing the frequency and increasing the magnitude of the second order system's resonant peak (Becker and Mote, 1990). That fatigue study used two exercises to fatigue the finger: 1) repeated lifting of a 5.88 N (1.32 N) weight 75 times with the palmar interosseus muscle and then 75 times with the dorsal interosseus muscle, and 2) maintaining a maximum voluntary isometric contraction for one minute. Each of these exercises is much more strenuous than the short, intermittent squeezes required of subjects in the current experiment.

4.4.1 SUGGESTED 4TH ORDER MODEL

A second order model paradigm fails to explain the change in moment of inertia estimates for increasing grip force, assuming that grip posture remains constant. Creating a fourth order model by treating finger dynamics, fingerpad dynamics, and actuator dynamics separately provides an explanation for the changing moment of inertia estimates. Figure 13 treats the finger, fingerpad, and actuator as two cascaded second order systems. J_1 , B_1 , and K_1 represent the dynamics of the finger, connected to ground. B_2 and K_2 represent the damping and stiffness coefficients of the fingerpad, and J_2 represents actuator moment of inertia. If the finger were to become infinitely stiff, the finger mass would become rigidly connected to ground, and the driving force (torque) would have no effect on finger motion, preventing the inclusion of finger parameters in the estimate, which would become dominated by fingerpad and actuator dynamics. One would expect a second order estimate of this fourth order system to have a moment of inertia estimate dropping to equal just the moment of inertia of the actuator. Empirical results come close to this behavior, though the moment of inertia estimates actually drop *below* the reference estimate for actuator moment of inertia for the strongest grip forces,

implying that this model does not completely explain the experimental moment of inertia observations.




Figure 13: Suggested fourth order model, with finger dynamics, fingerpad dynamics, and actuator dynamics separately represented.

Figure 14 presents the mechanical network corresponding to the system in Figure 13, and Equation 1 presents the corresponding system equations. Using the state-space model of the system described in Equation 2, a simulation was conducted that confirmed that such a fourth order system being observed and analyzed with a second-order model would result in a moment of inertia estimate dropping to equal the actuator moment of inertia as finger stiffness and damping increased to be much higher than fingerpad stiffness and damping.




Figure 14: Fourth order mechanical network.

Equation 1: Network equations for fourth-order mechanical network.

$$T(t) = J_2 D^2 \theta_b + B_2 D(\theta_b - \theta_a) + K_2 (\theta_b - \theta_a);$$

$$0 = (J_1 D^2 + B_1 D + K_1) \theta_a - B_2 D(\theta_b - \theta_a) - K_2 (\theta_b - \theta_a),$$

Equation 2: State-space model for fourth-order system.

$$\begin{bmatrix} \dot{\theta}_b \\ \ddot{\theta}_b \\ \dot{\theta}_a \\ \ddot{\theta}_a \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -K_2 & -B_2 & K_2 & B_2 \\ J_2 & J_2 & J_2 & J_2 \\ 0 & 0 & 0 & 1 \\ K_2 & B_2 & -(K_1 + K_2) & -(B_1 - B_2) \\ J_1 & J_1 & J_1 & J_1 \end{bmatrix} \cdot \begin{bmatrix} \theta_b \\ \dot{\theta}_b \\ \theta_a \\ \dot{\theta}_a \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} \cdot T$$

For the simulation used to test this hypothesis, a second-order estimate of system dynamics was made using the least-squares fitting method presented earlier, operating on data from an example trial from Subject 5 (selected at random) at the highest grip force, and assuming a relationship between the stiffness and damping coefficients of this estimate and the stiffness and damping coefficients of the 4th order estimate defined by the equations:

$$K_T = \frac{K_1 K_2}{K_1 + K_2}; \quad K_1 = 10K_2; \quad B_T = \frac{B_1 B_2}{B_1 + B_2}; \quad B_1 = 10B_2;$$

where K_T and B_T are the second order fitted model parameters. In other words, the two springs were treated as series springs with the finger stiffness, K_1 , much greater than the fingerpad stiffness, K_2 , and the finger damping coefficient was estimated to be much greater to the fingerpad damping coefficient in a similar fashion. The fourth order model assumes a moment of inertia for the motor, J_2 , of 19.8 gm-cm² (derived in Appendix C) and assumes a generous estimate of 3 gm-cm² for the finger moment of inertia, J_1 . These estimates are not intended to be definitive estimates of actual system parameters. They represent hypothetical assertions to see if there exists a set of fourth-order physical parameters that could produce the behavior observed.

Figure 15 shows the displacement response to a perturbing pulse by the estimated fourth order model (dashed lines), the estimated second order model (solid line), and the original measured response. The fourth-order response appears very close to both the second-order model response and the original measured response. Figure 16 provides more direct support to the hypothesis that very high finger stiffnesses can cause fourth-order behavior that reduces the moment of inertia apparent to the current second-order estimation technique. It shows that the acceleration profile of the knob/fingerpad node perfectly matches the second-order acceleration profile (the dashed line has been offset upward slightly for ease of viewing; otherwise the two curves would overlap perfectly). The finger acceleration (available only in the fourth order model) is much lower, and the effects of its motion are not detectable by observing the knob/fingerpad node. Since observed moment of inertia is directly proportional to acceleration, this means that when the finger stiffness greatly exceeds the fingerpad stiffness, it can effectively hide the finger moment of inertia from the second-order estimation technique employed for these experiments.

Figure 15 shows one obvious difference between the two estimated responses and the measured response that appeared in nearly all the responses for the two highest grip forces for nearly all subjects; an initial peak, followed by a dip below the modeled response, then a rise. None of the linear models presented here account for that behavior. The consistency with which the phenomenon appears and the failure of the linear models presented so far to deal with it clearly indicates an unmodeled system characteristic; however, since the phenomenon appears only for the highest grip forces, which are

higher than occur in realistic use of haptic knobs, this shortcoming does not threaten the practical utility of the models presented.

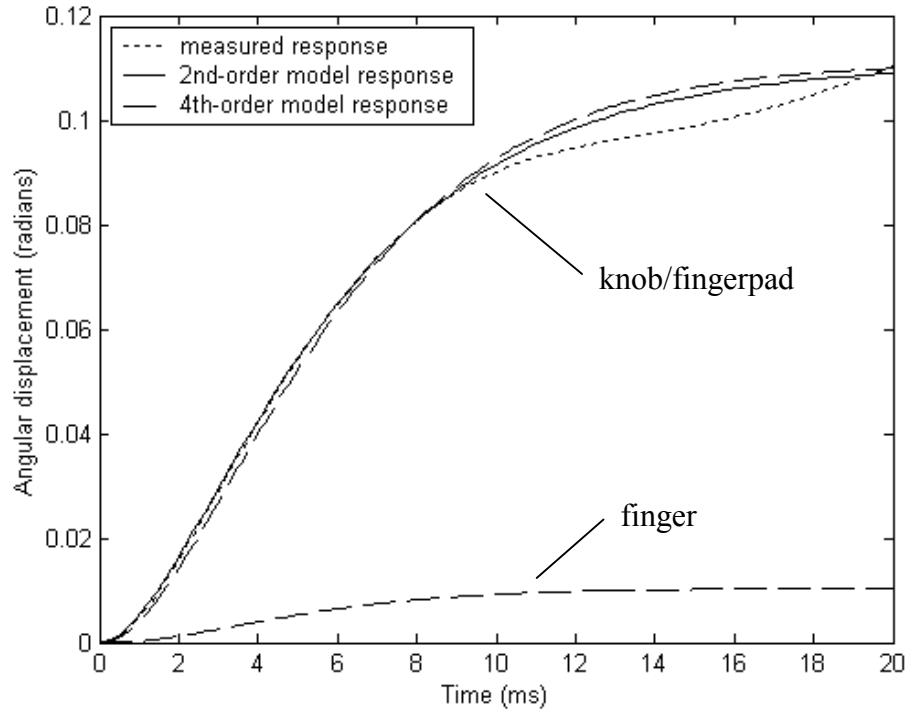


Figure 15: Displacement Responses of Estimated 4th and 2nd-Order Systems Compared to Measured Response

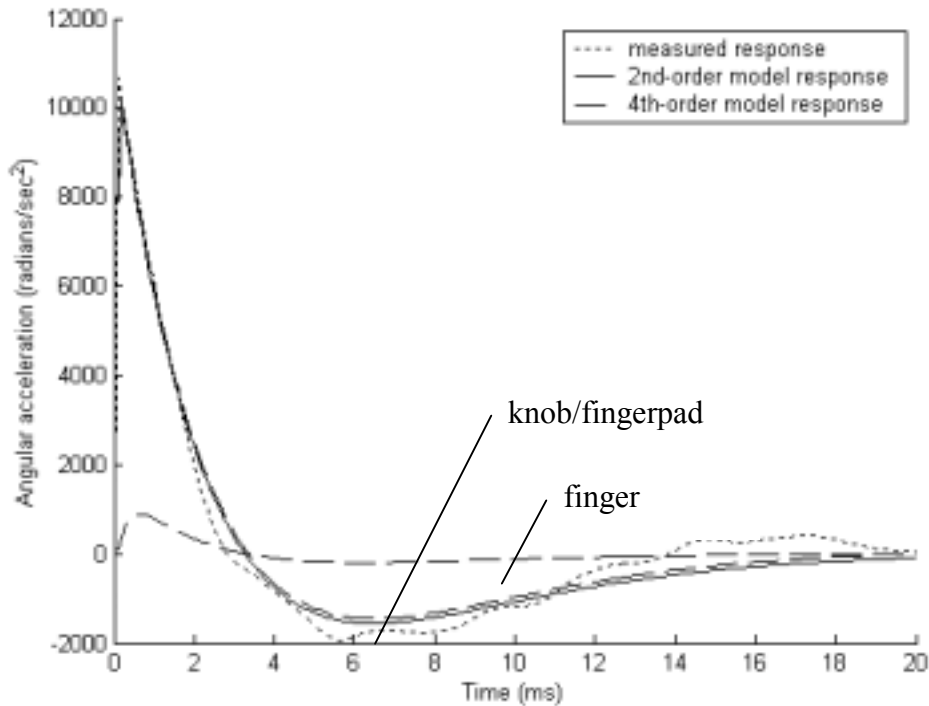


Figure 16: Acceleration Responses of Estimated 4th and 2nd-Order Systems Compared to Measured Response

Karason and Srinivasan (1995) had similar difficulties fitting a second-order model to the dynamics of two fingers grasping a linear active instrumented object that could expand or contract suddenly (one could describe the HREF haptic knob as a rotational active instrumented object). They successfully identified a third-order linear model that treats the impedance of the finger and the fingerpad separately. Their third-order linear model is a close analog of the fourth order rotational model proposed here (which would be third order if the actuator mass were subtracted). Karason and Srinivasan consider finger flexion/extension, and the fingerpad in compression, while the present work considers grasp rotation involving primarily abduction/adduction, and the fingerpad in shear.

4.5 Conclusions

Results show that for light to moderate grip forces, a second order linear lumped parameter model provides an excellent estimate of the dynamics of a human hand grasping a knob. Both intersubject and intrasubject variability are nontrivial, but compatible with reliable estimates of human grasp dynamics. Stronger grip forces challenged the second order model assumption. A higher order model that treats finger impedance and fingerpad impedance separately provides a substantially better explanation of combined system behavior. The system models provided here should provide a useful starting point for simulations of haptic knob systems grasped by hands.

Chapter 5 : System Identification for More Grasps and Conditions

5.1 Introduction

The wide variability in grasp dynamics exhibited in the previous investigation of a single grasp motivated further investigation of dynamics for different grasp postures and different-sized haptic knobs. Chapter 4 presented results for subjects grasping a 17.8 mm knob using a simple pinch grasp. This chapter explores the grasp dynamics for four more situations: a key pinch grasp of a 17.8 mm diameter knob, a simple pinch of a 29.9 diameter mm knob, a key pinch of a 29.9 mm knob, and a three-fingered grasp of a 29.9 mm knob.

5.2 Methods

Four more experiments were conducted, each with a different knob size and/or grasp posture, but following the same experimental procedure established for the first experiment presented in the previous chapter. The procedure involved one training block and three experimental blocks of trials for each experiment. Subjects ran the experiment with a key pinch grasp on the 17.8 mm knob on one day, and three experiments for the 29.9 mm knob on another day. The testing sessions with three experiments each lasted

approximately 45 minutes, of which approximately 1/3 of the time was spent in intermittent rest periods in between blocks. Five subjects participated in the experiment, three males and two females, all of whom participated in the first experiment.

Figure 17 presents the four grasp postures. Figure 17(a) shows a key pinch grasp on the 17.8 mm diameter knob, with the subject grasping the knob between the thumb and the side of the distal phalanx of the index finger. Figure 17(b) shows a simple pinch grasp of the 29.9 mm diameter knob, with the index finger and thumb aligned directly along the axis of the grip force load cell. Figure 17(c) shows a key pinch grasp of the 29.9 mm knob. Figure 17(d) shows a three-fingered grasp of the 29.9 mm knob, with the index and middle fingers touching each other and directly opposing the thumb.




Figure 17: Photographs of four knob grasps. (a) key pinch on 17.8 mm knob; (b) simple pinch on 29.9 mm knob; (c) key pinch on 29.9 mm knob; (d) three-fingered pinch on 29.9 mm knob.

5.3 Results and Discussion

Figure 18 shows results compared for both Experiment 1, a simple pinch grasp of the 17.8 mm knob (presented in the previous chapter), and Experiment 2, a key pinch of the 17.8 mm knob. Moment of inertia estimates are nearly identical for medium to high grip forces, and slightly lower for low grip forces though that difference lies within the variation observed in both cases. Damping is consistently higher for the key pinch grasp, possibly due to the larger amount of tissue in contact with and near the knob to dissipate energy with the key pinch grasp. Stiffness starts out identical for both cases, and then steadily becomes greater for the key pinch grasp. The relationship between stiffness and damping for the two cases causes the damping ratio to be higher for low grip forces with the key pinch grasp than with the simple pinch grasp, making the key pinch a preferred grasp to avoid limit cycles when contacting a virtual wall, and the simple pinch preferred for exciting limit cycles (as when a haptics expert anecdotally explores system robustness). Though direct comparison is difficult, it is interesting to note that twisting the knob with a simple pinch grasp involves abduction/adduction of both the forefinger and the thumb, while the key pinch grasp involves flexion/extension of the forefinger and abduction/adduction of the thumb.

In his study of finger dynamics during flexion and abduction, Hajian found that fingers exhibited substantially more damping in flexion than in abduction (which tends to agree with the current results), though he found higher stiffness for abduction than for extension (which tends to conflict with the current results) (Hajian, 1997). This difference could result from the fact that Hajian's experiment involved pressure on the more compliant fingerpad for extension and pressure on the less compliant side of the

finger for abduction, while the current experiment involved the opposite – pressure on the fingerpad for the abduction (simple pinch) case, and pressure on the side of the finger for the extension (key pinch) case. Though caution is warranted because Hajian’s disturbances caused compression on the finger while the disturbances of the current experiment caused shear, this analysis provides a potential explanation for the apparent difference in stiffness results between Hajian’s experiment and the current experiment.

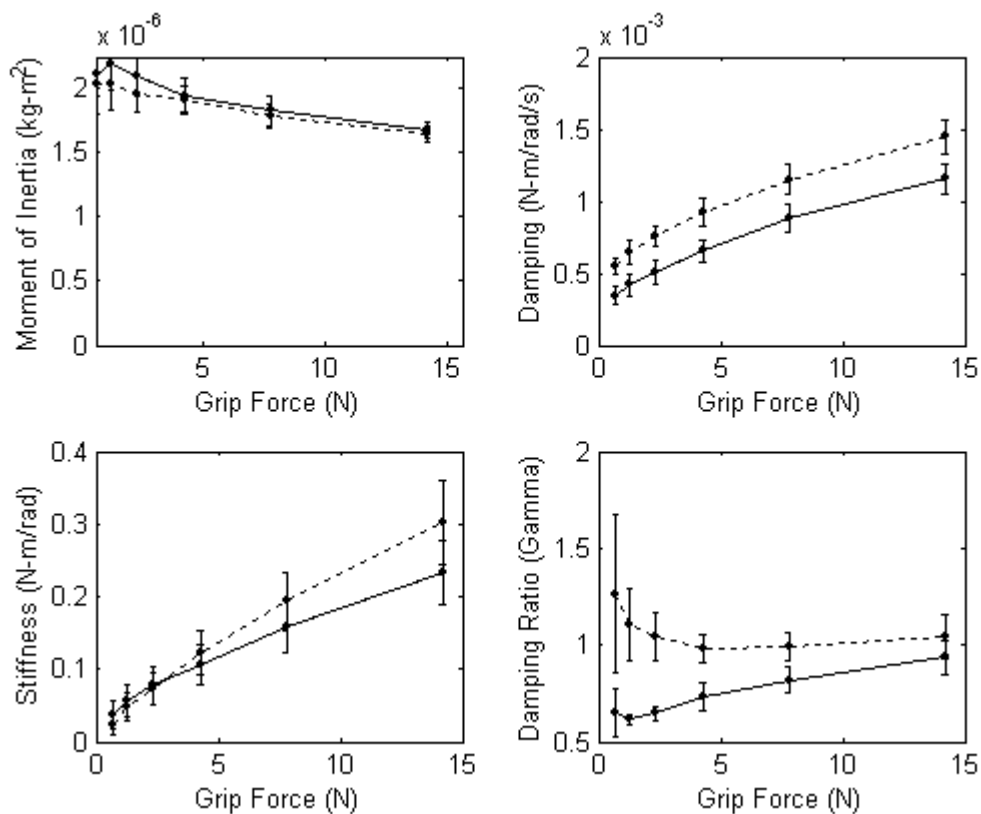


Figure 18: Comparison of Experiments for Different Grasps on the 17.8 mm Knob
Experiment 1, Simple Pinch Grasp (solid line)
Experiment 2, Key Pinch Grasp (dotted line)

Figure 19 shows results for Experiments 3, 4, and 5, for a simple pinch, key pinch, and three-fingered pinch, respectively, on the 29.9 mm knob. Moment of inertia estimates do not appear different. Detecting differences in moment of inertia was difficult with the 29.9 mm setup because the large inertia of the knob, $3.2 \times 10^{-6} \text{ kg}\cdot\text{m}^2$

(32.7 gm-cm², see Appendix C) caused the inertia of the fingers to be insignificant in comparison. The key pinch grasp exhibited consistently higher damping than the simple pinch grasp, a result very similar to that obtained for the same grasps with the smaller knob. The three-finger grasp exhibited even higher damping, not surprising given the additional amount of tissue “recruited” to dissipate energy with the involvement of an additional finger. Stiffness results for the simple pinch and key pinch cases are also comparable to the smaller knob experimental results, and adding a third finger increases stiffness significantly. Note the relatively large error bars, reflecting relatively high intersubject variability most likely contributed to by slight differences in grasp posture between subjects. The relatively higher damping than stiffness with the key pinch grasp again causes it to have a higher damping ratio for low grip forces than the simple pinch does, and also than the three-finger pinch.




Figure 19: Comparison of Experiments for Different Grasps on the 29.9 mm Knob
Experiment 3, Simple Pinch Grasp (solid line)
Experiment 4, Key Pinch Grasp (dotted line)
Experiment 5, Three-Finger Grasp (dot-dashed line)

Figure 20 shows results for all five experiments. Note the consistency of the moment of inertia results for the two knob sizes, and the steadily increasing damping and stiffness with the progression from Experiment 1 through Experiment 5. One would expect the larger knob experiments (3 through 5) to have higher rotational stiffness and damping coefficients, since the hand for a similar posture likely exerts similar forces on the different-sized knobs, but at a greater lever arm for the larger knob cases. A comparison of results normalized for knob radius, to create equivalent translational models along the knob circumference, would allow more insight into potential differences in dynamics between grasp postures with different knob sizes. Figure 21

contains such a comparison. To obtain the results in Figure 21, J, B, and K were divided by the square of the radius for a given knob to obtain translational values for the three parameters as they would be measured at the point the fingertips contact the knob surface.

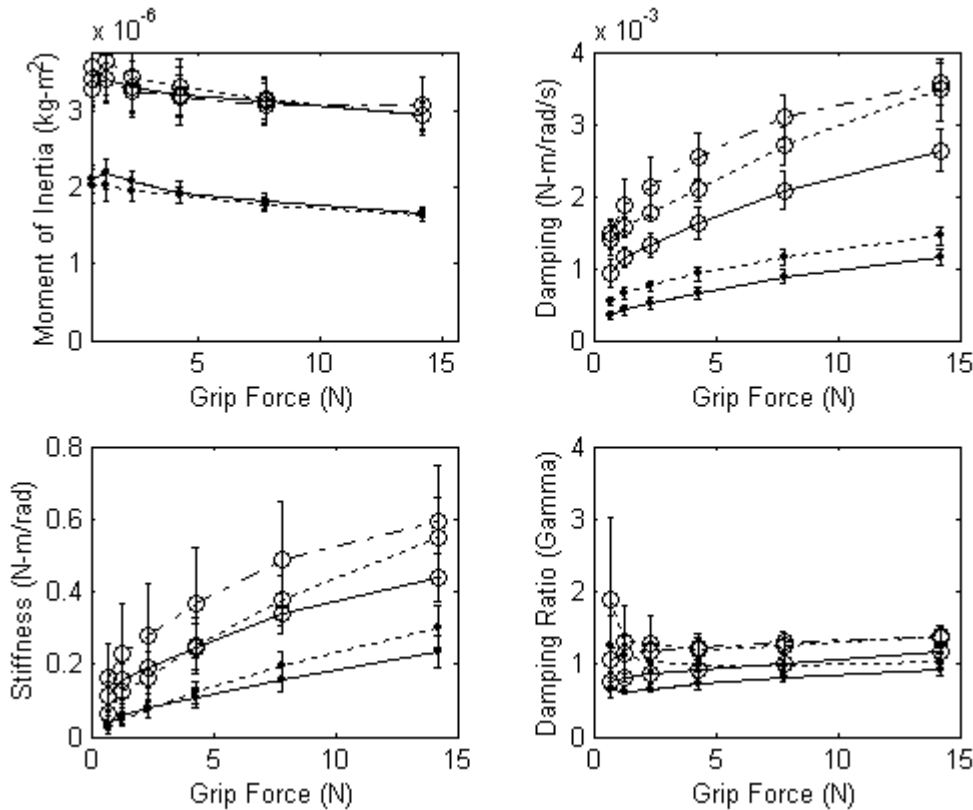


Figure 20: Comparison of Experiments for All Five Experiments
Experiment 1, Simple Pinch Grasp (solid line, with '●' data points)
Experiment 2, Key Pinch Grasp (dotted line, with '●' data points)
Experiment 3, Simple Pinch Grasp (solid line, with 'O' data points)
Experiment 4, Key Pinch Grasp (dotted line, with 'O' data points)
Experiment 5, Three-Finger Grasp (dot-dashed, with 'O' data points)

The moment of inertia plots for Figure 21 are not especially relevant, because they are dominated by the moment of inertia differences between the knobs themselves. The two simple pinch cases (solid lines) have lower damping than the other grasp postures. Both grasp postures for the small knob tend to have higher damping and stiffness than the

comparable grasp postures for the larger knob. A ready causal explanation does not present itself for this difference, though the grasps for the two different knob sizes involved significantly different flexion/extension postures for the thumb and forefinger, with numerous changes in kinematics, muscle and tendon loading, etc.

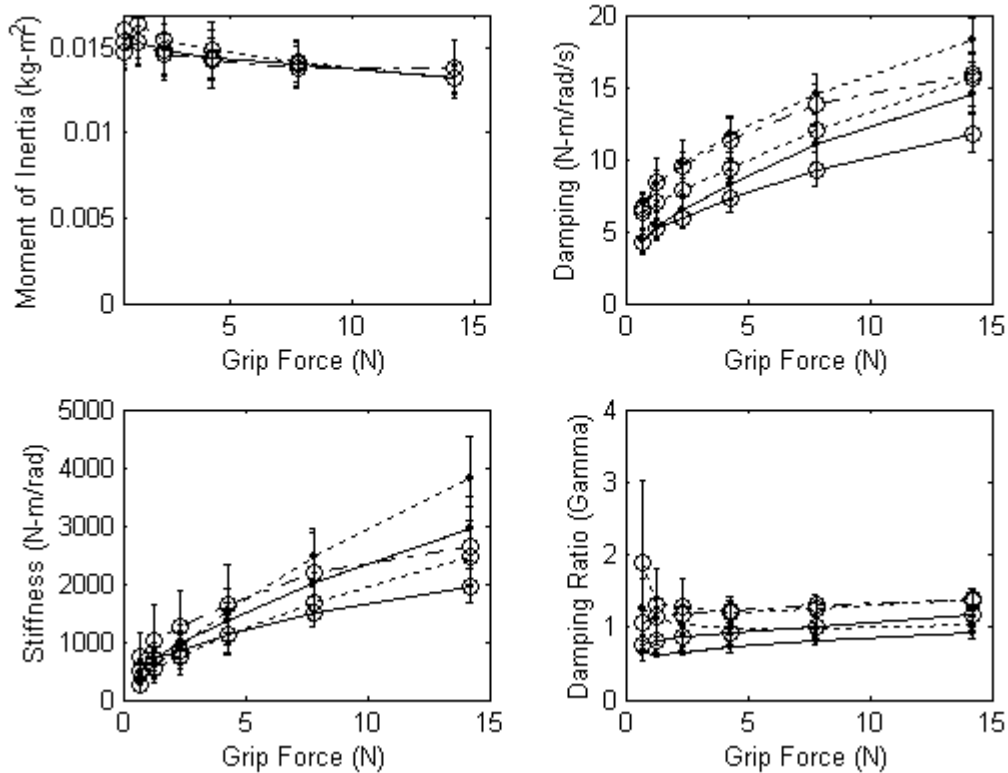


Figure 21: Comparison of Experiments for All Five Experiments, converted to NORMALIZED TRANSLATIONAL MODELS

- Experiment 1, Simple Pinch Grasp (solid line, with '●' data points)**
- Experiment 2, Key Pinch Grasp (dotted line, with '●' data points)**
- Experiment 3, Simple Pinch Grasp (solid line, with 'O' data points)**
- Experiment 4, Key Pinch Grasp (dotted line, with 'O' data points)**
- Experiment 5, Three-Finger Grasp (dot-dashed, with 'O' data points)**

5.4 Conclusions

Results show that across grasp postures and knob sizes, stiffness and damping increases with the involvement of more fingers in a grasp and with higher grasp forces.

A normalized set of translational models shows differences between the dynamics of

similar grasps with different knob sizes, potentially attributable to differences in kinematics and muscle and tendon loading between grasp postures for different knob sizes. The relatively large moment of inertia of the apparatus made lumped parameter estimates of the moments of inertia of the fingers difficult to obtain. Though the results presented in this chapter and the previous chapter do not allow one to exactly predict the dynamics of an arbitrary combination of grasp, knob size, grip force, and human subject, they do provide valuable models that researchers can use for studying control algorithms and that designers can use for estimating potential grasp dynamics that a haptic knob might experience.

Table 2: Results for Experiment 2, Key Pinch Grasp on 17.8 mm Knob, Across Subjects

Grip Force (N)	J (kg-m ²)			B (N-m/rad/s)			K (N-m/rad)			ζ (zeta)		
	J	std J	CV J	B	std B	CV B	K	std K	CV K	zeta	std zeta	CV zeta
0.71	2.02E-06	2.28E-07	11.3%	5.53E-04	5.44E-05	9.8%	0.02	0.01	61.0%	1.3	0.41	32.0%
1.3	2.03E-06	2.15E-07	10.6%	6.51E-04	8.34E-05	12.8%	0.05	0.02	40.7%	1.1	0.19	16.9%
2.4	1.96E-06	1.41E-07	7.2%	7.62E-04	6.94E-05	9.1%	0.07	0.02	30.3%	1.0	0.13	12.0%
4.3	1.90E-06	1.10E-07	5.8%	9.28E-04	9.63E-05	10.4%	0.12	0.03	24.3%	1.0	0.07	7.5%
7.8	1.78E-06	9.03E-08	5.1%	1.15E-03	1.05E-04	9.2%	0.19	0.04	20.6%	1.0	0.07	7.3%
14.2	1.64E-06	6.76E-08	4.1%	1.44E-03	1.21E-04	8.4%	0.30	0.06	19.3%	1.0	0.12	11.1%

Table 3: Results for Experiment 3, Simple Pinch Grasp on 29.9 mm Knob, Across Subjects

Grip Force (N)	J (kg-m ²)			B (N-m/rad/s)			K (N-m/rad)			ζ (zeta)		
	J	std J	CV J	B	std B	CV B	K	std K	CV K	zeta	std zeta	CV zeta
0.71	3.58E-06	4.45E-07	12.4%	9.36E-04	1.87E-04	20.0%	0.11	0.04	37.3%	0.8	0.06	7.8%
1.3	3.43E-06	2.12E-07	6.2%	1.16E-03	1.46E-04	12.6%	0.16	0.05	30.2%	0.8	0.06	7.4%
2.4	3.32E-06	3.34E-07	10.0%	1.32E-03	1.64E-04	12.4%	0.19	0.07	36.3%	0.9	0.13	15.4%
4.3	3.20E-06	3.88E-07	12.1%	1.63E-03	2.17E-04	13.3%	0.25	0.06	25.5%	0.9	0.11	12.2%
7.8	3.15E-06	3.06E-07	9.7%	2.08E-03	2.59E-04	12.5%	0.34	0.06	16.5%	1.0	0.15	14.4%
14.2	2.94E-06	1.76E-07	6.0%	2.63E-03	2.91E-04	11.1%	0.44	0.07	15.6%	1.2	0.15	12.9%

Table 4: Results for Experiment 4, Key Pinch Grasp on 29.9 mm Knob, Across Subjects

Grip Force (N)	J (kg-m ²)			B (N-m/rad/s)			K (N-m/rad)			ζ (zeta)		
	J	std J	CV J	B	std B	CV B	K	std K	CV K	zeta	std zeta	CV zeta
0.71	3.42E-06	3.69E-07	10.8%	1.41E-03	2.34E-04	16.6%	0.06	0.03	51.3%	1.9	1.13	59.6%
1.3	3.65E-06	5.41E-07	14.8%	1.56E-03	1.43E-04	9.2%	0.12	0.06	45.0%	1.3	0.51	38.6%
2.4	3.44E-06	5.25E-07	15.2%	1.76E-03	1.19E-04	6.8%	0.16	0.07	42.8%	1.3	0.39	30.0%
4.3	3.32E-06	3.77E-07	11.4%	2.09E-03	1.48E-04	7.1%	0.25	0.08	31.0%	1.2	0.22	18.1%
7.8	3.17E-06	2.62E-07	8.3%	2.70E-03	2.58E-04	9.6%	0.38	0.07	18.5%	1.3	0.09	7.3%
14.2	2.95E-06	1.95E-07	6.6%	3.48E-03	4.25E-04	12.2%	0.55	0.11	19.8%	1.4	0.16	11.3%

Table 5: Results for Experiment 5, Three-Finger Grasp on 29.9 mm Knob, Across Subjects

Grip Force (N)	J (kg-m ²)			B (N-m/rad/s)			K (N-m/rad)			ζ (zeta)		
	J	std J	CV J	B	std B	CV B	K	std K	CV K	zeta	std zeta	CV zeta
0.71	3.28E-06	2.01E-07	6.1%	1.48E-03	2.07E-04	14.0%	0.17	0.09	55.2%	1.1	0.16	15.4%
1.3	3.44E-06	3.09E-07	9.0%	1.88E-03	3.67E-04	19.5%	0.23	0.14	59.4%	1.2	0.24	20.7%
2.4	3.27E-06	2.67E-07	8.2%	2.12E-03	4.12E-04	19.4%	0.28	0.15	52.2%	1.2	0.15	12.5%
4.3	3.20E-06	2.79E-07	8.7%	2.54E-03	3.41E-04	13.4%	0.37	0.15	41.7%	1.2	0.16	12.8%
7.8	3.09E-06	2.76E-07	8.9%	3.09E-03	3.10E-04	10.0%	0.49	0.16	33.1%	1.3	0.14	10.9%
14.2	3.07E-06	3.89E-07	12.6%	3.56E-03	2.93E-04	8.3%	0.59	0.15	26.2%	1.4	0.17	12.4%

Chapter 6 : Simulation of Virtual Barrier Contact for Various Displacement Resolutions and Sample Rates

6.1 Introduction

Quantization in time due to discrete sampling, quantization of position information (e.g., from an optical encoder), and other quantization (e.g., rounding errors, output command quantization) can all degrade the performance of a haptic system by producing effects such as limit cycles upon contact with virtual barriers. Previous chapters identified a valid dynamic model and ranges for the dynamic parameters of a human hand in a knob grasp. This model and these parameters are prerequisites for simulations to explore the effect of displacement resolution and sample rate on the amplitude and frequency of limit cycles encountered upon contact with unilateral virtual barriers.

In systems where more than one contributor to limit cycles is present, it can be difficult to determine which one dominates for a given set of conditions. Identifying the prime cause of limit cycles within the desired region of operation allows one to focus redesign efforts to improve performance in the most efficient manner. This chapter includes an analysis of the relative contributions of sample rate and position quantization to limit cycle behavior for a range of parameters typical of a haptic knob grasped by a

user's hand. It assumes virtual wall force calculations based only on position, with no added damping, deferring an analysis of the effect of position quantization on the quality of derived velocity signals and the ramifications of these effects.

6.2 Previous Work

In the past 10 years, many investigators have examined haptic stability issues. Some have offered analysis methods to guarantee conditions for stable operation, or for operation without limit cycle oscillations. Others have offered mitigation or performance-enhancement techniques designed to reduce, rather than just identify, the stability limitations of haptic systems.

Colgate and Schenkel (1994) described necessary and sufficient conditions for the passivity of a class of sampled data systems (continuous-time plant and discrete-time controller) that includes linear constraints (e.g., haptic springs) and nonlinear unilateral constraints (i.e. haptic walls). They developed passivity conditions that reduce to:

$$b > \frac{KT}{2} + |B|,$$

where b is the inherent physical damping in the system, K is the virtual wall stiffness, T is the sample period, and B is the virtual damping. This finding shows that some physical dissipation must be present to enable passivity. Colgate and Schenkel also point out that given fixed virtual and physical damping, the maximum permissible stiffness is proportional to the sampling rate. Note that this inequality also permits *negative* virtual damping, raising the possibility of creating a system with significant amounts of inherent damping that could then be partially reversed by virtual negative damping (and almost completely reversed if $K=0$).

Passivity does not equate with stability. Passivity is a more stringent condition that forbids generation of energy. Miller and Colgate have worked to reduce the conservatism in passivity analyses (Colgate, 2001; Miller et al., 1999a, 1999b). While Colgate and Schenkel have shown that virtual damping can compromise passivity, this is often not a practical concern where sufficiently good displacement resolution and sample rates create good enough damping that the filtering effect of human and haptic device dynamics obscures differences between a virtual damper and a physical damper. Very high virtual damping values can cause difficulty, and poor quality velocity signals (e.g., those obtained from too heavily quantized displacement signals) severely limit the permissible virtual damping.

Gillespie and Cutkosky (1996) examine two ways in which time quantization contributes to chatter upon contact with a virtual wall, describing both as “energy leaks” that allow a supposedly passive virtual wall to transfer energy to the user which manifests itself in chatter. The first energy leak arises from the effects of the zero-order hold (ZOH) on the command output common to most sampled data systems. As the user penetrates the virtual wall, the ZOH produces a “staircase” output, underestimating the modeled wall force ($F_{wall} = -kx$) by an increasing amount as the time since the last sampling event increases. The output is only “correct” at the instant of sampling (neglecting computation delay) and is underestimating the modeled force at all other times. As the user withdraws from the virtual wall, the opposite situation occurs. At a sampling instant, the resultant output command exactly matches the virtual wall model, but as the manipulandum moves back towards the wall boundary and the virtual wall

force should lessen, the commanded output remains inappropriately strong until the next sampling instant.

This situation, with too little force exerted against the user over the distance of penetration, and too much force exerted against the user over the same distance upon withdrawal, results in the wall doing work on the user (delivering energy). Figure 22 reproduces Figure 1 from Gillespie and Cutkosky (1996). Commanded force appears as a solid line. While force for the actual simulation followed the equation $f_k = -Ky_k$, where K denotes wall stiffness and k is the sample index, for purposes of the illustration commanded force has been normalized by K and multiplied by -1 , so that it follows the equation $f_k = 1y_k$, effectively representing the quantization of the position signal and allowing for a direct comparison to the continuous position signal (dotted line). The staircase plot illustrates the underestimation of wall force during penetration and the overestimation of wall force during withdrawal that characterizes this ZOH-induced energy leak.

The second energy leak occurs due to asynchrony between the time the manipulandum penetrates the virtual wall boundary and the time that the sampled-data system recognizes this crossing and updates the control law to include wall stiffness in the calculation of the output command. A similar asynchrony occurs between the time the manipulandum exits the virtual wall boundary and the time that the system recognizes this crossing and updates the control law to remove wall stiffness from the calculation of the output command. As a result, the manipulandum momentarily experiences too little force upon wall entry and too much force upon wall exit, delivering energy to the user. Figure 22 shows the delay, Δt_a , between initial wall penetration and the first control law

update recognizing the penetration. The second delay, Δt_b , represents the delay between the time the manipulandum exits the wall and the time the controller turns off the wall force. Tsai and Colgate (1995) also acknowledged the energy surplus caused by the asynchrony upon wall exit.




Figure 22: Plot of modeled manipulandum position and control effort (from Gillespie and Cutkosky, 1996).

Ellis et al. also addressed the energy-generating aspects of ZOH, describing it as a particularly poor approximation of the desired continuous forcing function (Ellis et al., 1996, 1997). For situations where the force output occurs at the same rate as the force calculation (or sensor acquisition), Ellis et al. present an approximation scheme that uses ZOH hardware to effectively approximate a first-order hold response. By taking the current calculated force and the next anticipated force (determined by linear extrapolation from the current and previous forces), they estimate the desired ZOH as the average of the current and anticipated force outputs. For a perfect estimate of the next anticipated force, this scheme would apply just the right amount of force to the system to

approximate the continuous desired forcing function without instilling any excess energy. Ellis et al. acknowledge the error of the first-order hold approximation, and add a correcting term to account for force prediction error from the last prediction cycle.

In contrast to Gillespie's approach, Ellis' approach does not require (or does not take advantage of, depending on one's perspective) any dynamic model information for the user or haptic device. Since Ellis' scheme relies only on force output data and not sensor data (its lack of dependence on velocity information is especially notable), Ellis' approach could be well suited for application to inexpensive haptic systems with coarse-resolution displacement sensors. For situations where the force calculation (or sensor acquisition) can occur faster than the force output, Ellis suggests using the larger number of calculated forces to make a better approximation of the next anticipated force (such as by using a least-squares fit polynomial). For situations where the force output occurs faster than the force calculation (or sensor acquisition), Ellis suggests breaking up the force output into multiple ZOH steps intended to approximate the desired output force trajectory from the current force to the next computed force.

Love and Book (1995) used Jury stability criterion (Jury, 1964) to define a relationship between the target stiffness, target damping, and sample rate of a system, and found an upper bound on the permissible stiffness of a wall. They also found a lower and upper bound on the target damping of the system. Love and Book did not account for the dynamics of the human.

Adams et al. described a two-port model from network (circuit) theory for the design of unconditionally stable haptic interfaces (Adams and Hannaford, 1998, 1999, 2001; Adams et al., 1998, 2000). Their approach uses a virtual coupling scheme which is a

general implementation preceded by similar efforts for particular cases by Colgate et al. (1995) and Zilles and Salisbury (1995). Colgate et al. (1993a) first proposed using a virtual coupling between a haptic display and a virtual environment, though the idea has precedent in the telerobotics literature (Anderson and Spong, 1992; Colgate et al., 1993b; Hannaford, 1989). Adams and Hannaford use Llewellyn stability criteria to design virtual couplings for both impedance control and admittance control cases that guarantee stability, assuming a passive virtual environment and a passive human user. A typical virtual coupling for impedance control would be composed of a parallel spring and damper inserted between the haptic display and the virtual environment. The general approach works for any of the four possible combinations of impedance or admittance display and virtual environments. Adams and Hannaford account for the ZOH and design virtual couplings to guarantee stability despite it, but do not appear to directly mitigate the effects of the ZOH to increase system performance. For passive virtual environments, these approaches do have the tremendous benefit of separating haptic display control system design from virtual environment design.

Miller et al. (1999a) extended virtual coupling theory to define conditions guaranteeing the absence of oscillations for a class of nonlinear virtual environments while permitting those environments to be non-passive. Further work by Miller et al. (1999b) defines oscillation-free conditions for linear virtual environments that may be non-passive. Both approaches rely on “excess passivity” (damping) in the operator/haptic device subsystem to create non-oscillatory conditions even in the presence of a non-passive virtual environment. These analysis techniques would not

necessarily allow one to extend the performance of haptic systems with low damping, such as most commercial haptics systems.

Tsai and Colgate (1995) took advantage of particular properties of a unilateral nonlinear constraint (a typical haptic wall with no force outside of a boundary, and a spring force within the wall boundary) to define criteria on a discrete-equivalent transfer function for the combined operator/haptic device/environment system to guarantee the absence of oscillations. Their technique uses properties of the nonlinearity to allow them to be less conservative than Tsytkin's Criterion (1962), and they are able to graphically represent their constraint in the Nyquist plane, allowing comparison to other methods such as Tsytkin's that can be interpreted in the Nyquist plane. In order to apply the analysis, one must have knowledge of the haptic system dynamics as well as the human operator dynamics (with an assumption that the operator is a linear time-invariant system).

Hannaford and Ryu (2001a, 2001b) implemented a passivity observer that measures energy flow into or out of subsystems, and a passivity controller that adaptively introduces a dissipative term that absorbs any excess energy produced by the system. They cite precedence for passivity concepts in the area of force-feedback teleoperation with time delay (Anderson and Spong, 1992; Niemeyer and Slotine, 1991), but state that traditional passivity approaches are more conservative than necessary. One example of a less conservative approach is the two-port approach of Adams and Hannaford, cited above.

Hannaford and Ryu achieved good results with their passivity controller, quickly damping out unwanted oscillations for contact with high-stiffness virtual walls and for

situations with time delay up to 15 ms. They cite two limitations of their approach. First, the passivity controller may call for forces that exceed actuator limitations. In this situation, the passivity controller might take longer to dissipate the accumulated energy. Second, the approach relies on a velocity signal to compute system energy flow, and noise in a derived velocity signal can reduce the feasible gains of the passivity controller, again limiting the speed with which it can dissipate energy. The latter constraint is especially salient for systems with low-resolution optical encoders, where derivation of a good velocity signal is particularly difficult.

In published work, Hannaford and Ryu (2001a, 2001b) suggest that knowledge of the human damping parameter would allow the passivity controller to be set to permit some positive energy flow out of the haptic device (which could then be dissipated by the human). This would mean that the passivity controller could invoke itself less often, perhaps leading to less distortion in the intended command signal; however, in subsequent unpublished experiments, they found that the passivity controller was invoked so seldom in any case that loosening energy constraints to allow it to be invoked less would have no practical benefit (Hannaford, 2001). These findings were obtained with a large, high-inertia haptic device; lower inertia devices might yield different results. The conclusion that knowledge of the human damping is unnecessary, combined with the fact that knowledge of human stiffness and inertia are also unnecessary for the technique, mean that Hannaford and Ryu's passivity controller can enhance system stability in a less conservative fashion than two-port coupling methods, without any knowledge of human system dynamics.

6.3 Does Position Quantization Possess Inherent Energy Leaks?

Position quantization such as that present with optical encoders has been observed to worsen performance of haptic systems simulating virtual environment components such as virtual walls. Understanding the role of position quantization in performance degradation is one of the goals of this work. Gillespie (1996) expressed the deleterious effects of time quantization in terms of an "energy leak," showing that the zero-order hold produced staircase plots (see Figure 22) illustrating that each bounce of the manipulandum against the virtual wall results in more energy being delivered to the human. One might expect displacement quantization to produce a similar effect, but this is not the case.

Figure 23 illustrates a simulation of a hypothetical continuous-time system with virtual wall spring forces created from a quantized displacement signal. The stepped line indicates commanded controller effort (scaled and adjusted for plotting in a fashion similar to the command trajectory in Figure 22). The other line indicates manipulandum displacement. Ball displacement was calculated by sequential invocations of an ordinary differential equation (ODE) solver in Matlab (ode23) which uses an explicit Runge-Kutta (2,3) pair (Bogacki and Shampine, 1989; The Mathworks, 2001a,b). The ODE solver was invoked once for each period the manipulandum was outside the wall, and once for each encoder interval within the wall. A state machine in the invoking program and an event detector in the ODE solver stopped the solver at the threshold of each new encoder value, updated the computed wall force with zero latency as soon as the encoder value incremented or decremented, and called the ODE solver again

This simulation has no damping, so energy should ideally be neither gained nor lost. Figure 23 contains three peaks of equal value, and demonstrates a constant energy state. Inspection of the staircase plot reveals quantization of displacement, with updates as soon as a new displacement quantum is reached, and a *symmetrical*, rather than a lagging staircase plot. With this staircase, the user pushes a little harder than he or she would with an ideal system while penetrating the wall, and the wall overestimates forces by the same amount during withdrawal; the user does the same amount of work on the wall during penetration as the wall does on the user during withdrawal. There is no net energy gain. This simulation demonstrates that displacement quantization alone cannot cause energy leaks of the type observed by Gillespie (1996).

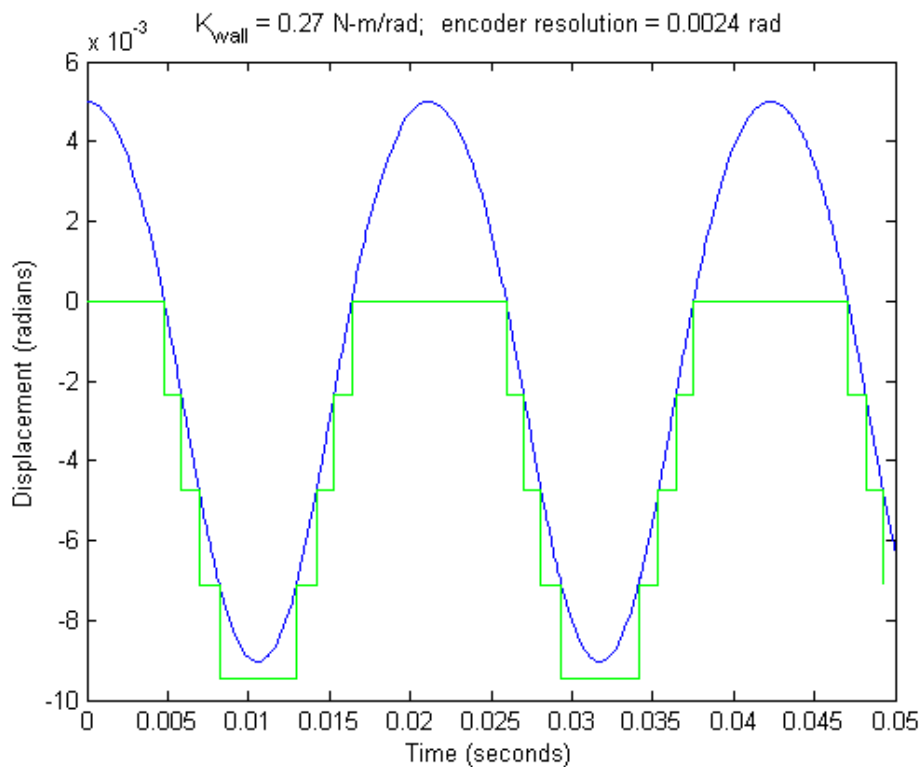


Figure 23: Continuous-Time Simulation with Encoder Displacement Quantization

A reader searching for potential hazards of low resolution displacement signals might imagine a situation where a user approaches a wall very slowly, and a highly

quantized position signal feeding the virtual wall control law $f_k = -Ky_k$ causes an extremely high force upon first penetration of the virtual wall, making the manipulandum rebound violently, and persistently adding energy to the system with each bounce. Figure 24 shows that this hypothetical scenario does not occur for a highly quantized displacement signal in a continuous time system. The three plots represent three different simulated trials of a continuous-time system with encoder quantization, and various initial conditions for displacement (again, with no physical damping). From top to bottom in the figure, the trials possess decreasing magnitudes for the displacement initial condition, with each successive trial representing a smaller initial energy state.

The continuous-time nature of this system means that the wall force turns on at the precise moment the manipulandum enters the wall, and turns off at the precise moment it exits the wall. As for the previous continuous-time encoder simulation, work done by the manipulandum on the wall during penetration equals work done by the wall on the manipulandum during withdrawal, and energy remains constant. The succession of plots in Figure 24 shows that the force pulses look a bit like pulse-width modulation pulses (with infinite time resolution) precisely titrated to deliver just enough energy to reverse the manipulandum and send it away from the wall with the same amount of energy it entered with. Though it appears that the manipulandum does not actually penetrate the wall in each case, this illusion is caused by the high wall stiffness. Figure 25 contains a closeup of one of the bounces clearly showing wall penetration. The presence of only one force step level for these examples is simply a special case; for larger initial energy states, less stiff walls, or higher-resolution displacement signals, the force trajectories would have a staircase pattern similar to that in Figure 23.




Figure 24: Continuous-Time Simulations of a Very Stiff Wall with Displacement Quantization

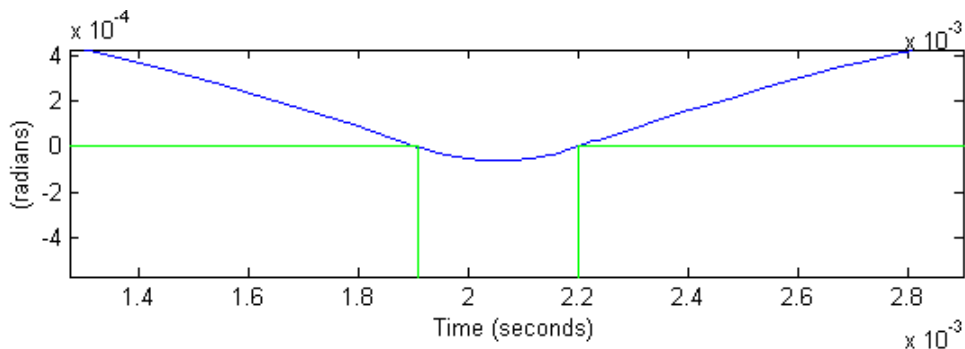


Figure 25: Closeup of a Bounce During Continuous-Time Simulation of a Very Stiff Wall with Displacement Quantization

Adding physical damping, even the small amount of damping typically present in a gentle grasp, rapidly quenches displacement oscillations. Figure 26 shows a simulation for exactly the same parameters as the simulation represented by Figure 23, but with the addition of 3.45×10^{-4} N-m/rad/s of damping; the average damping across all subjects for the gentlest grip in the grip dynamics identification experiment results in Table 1 on

Page 33. With this small amount of physical damping present, the oscillations of initial contact settle almost completely in 50 milliseconds.

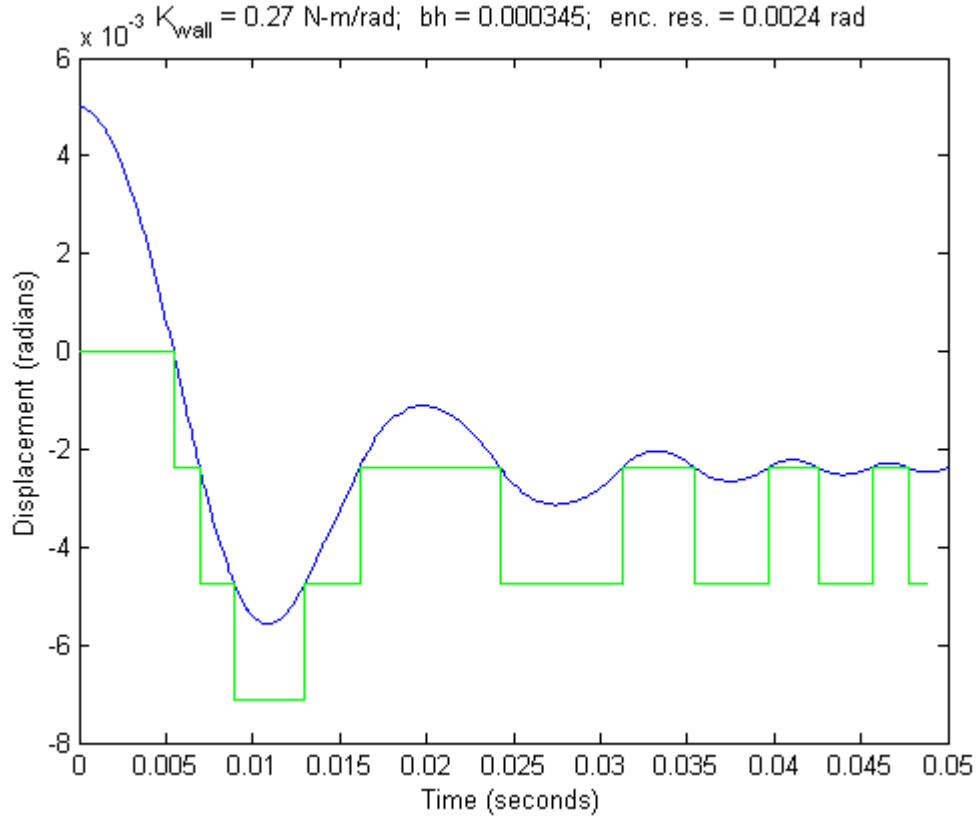


Figure 26: Continuous-Time Simulation with Encoder Displacement Quantization and Physical Damping

If displacement quantization does not have inherent energy leaks, then how can we explain the worsening of performance seen with lower resolution sensors? Though low resolution displacement sensors do not possess their own energy leaks, the following analysis will show that coarser resolution displacement sensors exacerbate energy leaks caused by the zero-order hold, and the worse the resolution, the worse this effect becomes. Numerous investigators have also commented on the difficulty of obtaining satisfactory velocity signals from low resolution displacement sensors, which further explains the burdens they place on system performance.

6.4 “Energy Leak” Analysis for a Haptic Knob

In order to analyze the effects of the zero-order hold (ZOH) and position quantization on stability, Gillespie’s simulations were extended to include position quantization as well as ZOH effects for various sampling periods. Figure 27 shows Gillespie’s models of a finger/manipulandum contacting a virtual wall, with one model for the “outside wall” case and the other for the “inside wall” case. Lowercase letters m , b , and k represent the mass, damping, and stiffness parameters, respectively, of the human finger. M represents the manipulandum mass, B represents the manipulandum damping, and K represents the damping of the virtual wall. g represents the bias force (supplied by gravity in the case of a “bouncing ball” simulation). In the present simulations, j and J will be substituted for m and M to represent the moment of inertia of the fingers in a knob grasp and the moment of inertia of the manipulandum (knob, motor, etc.), respectively. Unless otherwise indicated, g will be considered a constant bias torque imposed by a human user grasping the knob and twisting to maintain contact with a virtual barrier (a rotational “wall”).

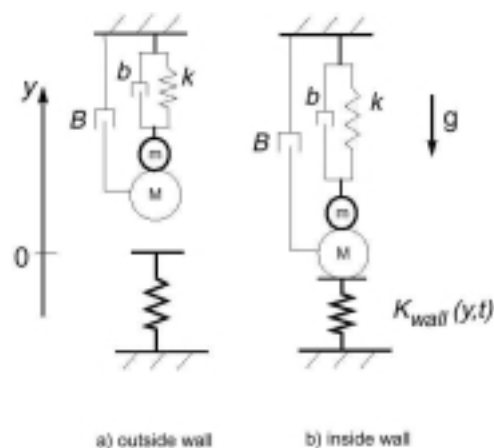


Figure 27: Gillespie's Model of a Finger/Manipulandum Contacting a Virtual Wall (from Gillespie, 1996)

6.4.1 INITIAL SIMULATIONS WITH INERTIA, NO DAMPING OR HAND STIFFNESS (THE “BOUNCING BALL” MODEL)

Simulations such as that illustrated in Figure 28 were performed for all combinations of seven sampling rates (100, 215, 464, 1000, 2154, 4642, and 10000 Hz) and seven encoder resolutions from 2^7 counts/revolution to 2^{13} counts/revolution (128, 256, 512, 1024, 2048, 4096, and 8,192 counts/revolution). Table 6 contains the rest of the simulation parameters. Ellis et al. (1996, 1997) showed that errors in estimating the wall force (from sources such as the ZOH approximation) will magnify the system energy quadratically. Figure 28 demonstrates this quadratic growth. A second-order polynomial fit to the bounce peaks was obtained for an estimate of energy growth during each trial. The second-order coefficient, loosely termed the “energy growth rate,” in units of J/s^2 , was recorded for later analysis.

Table 6: Simulation Parameters

Parameter	Symbol	Value	Units
Bias force	g	0.001	N-m
Wall stiffness	K	0.27	N-m/rad
Grasp stiffness	k	0	N-m/rad
Grasp damping	b	0	N-m/rad/s
Grasp inertia	j	3×10^{-7}	Kg-m ²
Knob/motor damping	B	0	N-m/rad/s
Knob/motor mass	J	1.9×10^{-6}	Kg-m ²




Figure 28: "Bouncing Ball" Simulation of Manipulandum Contacting Virtual Wall with Time Discretization and Displacement Quantization, but no Physical Damping. Sample Rate = 1 kHz; Encoder Resolution = 4096 counts/revolution.

Figure 29 contains a plot of energy growth rates as a function of sampling rate and encoder resolution (counts/revolution). Note that for the highest sampling rates, the growth rate remains near zero regardless of encoder resolution, illustrating the observation made earlier that encoder quantization alone does not contribute to energy leaks. For the best encoder resolution (8,192 counts/revolution), growth rate still increases with increasing sampling period, a situation close to that examined by Gillespie (who examined time quantization without position quantization). Figure 29 shows that over the full range of sampling rates and encoder resolutions, encoder resolution does not appear to have a significant effect on the energy growth rate. Figure 30 shows the same data as Figure 29, with the growth rate magnitude on a log scale to better represent the

smaller growth rates at high servo rates. The nearly planar mesh plot in Figure 30 shows that growth rate increases logarithmically with decreasing sample rate. Since the asynchrony between actual wall crossing and control law updates that happen once every sample period is a stochastic leak, the data in Figure 29 and Figure 30 represent an average of several simulations for each case, with an initial simulation step a different fraction of the standard sampling interval for that case (so that wall crossings happen at varying points in the simulation intervals for different simulation trials with the same set of conditions).




Figure 29: Quadratic Energy Growth Rates as a Function of Sampling Period and Encoder Resolution




Figure 30: Quadratic Energy Growth Rates as a Function of Sampling Period and Encoder Resolution (Log Magnitude for Growth Rate)

6.4.2 WITH HAND DYNAMICS ADDED

Figure 31 shows oscillations of a simulated virtual knob being gently held against a virtual barrier with hand dynamics representing the lightest grasp investigated in preceding chapters. Table 7 contains simulation parameters used to obtain the results in Figure 31. The sampling rate was 1000 Hz and the encoder resolution was 2650 counts/revolution. The damping and stiffness parameters for the lightest grasp measured in Chapter 4 were used. A light grasp with just enough bias torque into the virtual barrier to remain in contact with it represents the worst case for the generation of limit cycle oscillations. A light grasp has the lowest amount of damping available to dissipate energy and the lowest stiffness available to resist limit cycle excursions.

One might ask whether a haptic knob without anyone grasping it is actually the worst case, given that it has absolutely no stiffness and damping (neglecting potential damping in the haptic device). One might further ask why haptic designers should not design for this worst case and dispense with the need for system identification of user dynamics. In fact, no hand grasping the knob is not the worst case. With no hand grasping the knob, there is no bias torque, and the knob will be repelled from the barrier once and come to rest.

In contrast to the bouncing ball simulations with no dissipative elements presented above, a system with any physical damping will reach a steady-state oscillation magnitude. Injecting more energy into the system will result in faster oscillations, which will result in more dissipative loss through damping, creating a self-limiting balance.

Table 7: Simulation Parameters

Parameter	Symbol	Value	Units
Bias force	g	0.001	N-m
Wall stiffness	K	0.27	N-m/rad
Grasp stiffness	k	0.08	N-m/rad
Grasp damping	b	3.45×10^{-4}	N-m/rad/s
Grasp inertia	j	3×10^{-7}	Kg-m ²
Knob/motor damping	B	0	N-m/rad/s
Knob/motor mass	J	1.9×10^{-6}	Kg-m ²




Figure 31: Simulated Oscillations of a Lightly-Grasped Haptic Knob in Contact with a Stiff Virtual Barrier




Figure 32: Simulation of Hand Lightly Pressing Knob Against Stiff Virtual Wall, with Lines Fitted to Steady State Peaks and Troughs to Measure Limit Cycle Magnitude (2000 Hz, 8192 encoder counts/revolution)

To explore the effect of displacement and time quantization on haptic chatter, one would like to measure oscillation magnitude and frequency for a range of displacement and time quantization levels. Results will only be valid for specific hand dynamics and virtual barrier stiffness. Knowledge of how chatter varies with these parameters would also be interesting. Figure 32 shows a simulation similar to that in Figure 31, but for a higher sample rate and encoder resolution (note the higher bounce frequency and lower steady-state bounce magnitude). The simulation program monitors the bounces until reaching steady state and continues to run for ten more bounces, making linear estimates of average peak and trough values (see dashed lines in Figure 32) to determine limit cycle oscillation magnitude. Figure 33 through Figure 38 show the results of a series of

simulations using the parameters in Table 7, with sampling rate ranging from 100 Hz to 10,000 Hz, and encoder resolution ranging from 2^7 (128) counts/revolution to 2^{13} (8,192) counts/revolution.

Figure 33 through Figure 35 show the same data, with various combinations of logarithmic and linear scales for the three axes. While Figure 33 and Figure 34 give the most intuitive illustration of the sharp increase in limit cycle oscillation magnitude with worsening sample rate and encoder resolution, Figure 35 gives the most informative view, showing oscillation magnitude increasing in a steady logarithmic manner, both with worsening sample rate and with worsening encoder resolution. Figure 36 shows peak-to-peak oscillation magnitude, expressed in units of encoder counts. Figure 37 is a plot of the same data, saturated to emphasize that the peak-to-peak oscillations are less than +/- 1 encoder count for the majority of the cases.

Figure 38 shows oscillation frequency as a function of sampling rate and encoder counts/revolution. As sampling rate increases, oscillation frequency trends upward. Figure 39 displays the trend in oscillation frequency as a function of sample rate, averaged across different displacement resolution cases. Figure 40 shows that oscillation frequency trends only slightly downward with increasing displacement resolution.

Magnitude Plots:




Figure 33: Peak-to-Peak Oscillation Magnitude for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Period and Displacement Quantization




Figure 34: Peak-to-Peak Oscillation Magnitude for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Rate (log10) and Displacement Counts/Revolution (log2)




Figure 35: Peak-to-Peak Oscillation Magnitude (log10) for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Rate (log10) and Displacement Counts/Revolution (log2)




Figure 36: Peak-to-Peak Oscillation Magnitude, Expressed in Units of Encoder Counts




Figure 37: Saturated Mesh Showing Cases with Oscillation Magnitudes Greater than +/- 1 Encoder Count as "1" and Cases with Oscillations Less Than +/- 1 Encoder Count as "-1"

Frequency Plots:




Figure 38: Oscillation Frequency for Manipulandum in Contact with Stiff Barrier, with Physical Damping from Human Grasp, as a Function of Sampling Rate (log10) and Displacement Counts/Revolution (log2)




Figure 39: Average Limit Cycle Frequency as a Function of Servo Rate




Figure 40: Average Limit Cycle Frequency as a Function of Displacement Resolution

Chapter 7 : Describing Function Analysis

7.1 Introduction

Modern nonlinear control theory offers a method, called “describing function analysis,” to predict the magnitude and frequency of limit cycles. Describing function analysis creates quasi-linear frequency-domain approximations of nonlinearities. The method is particularly useful for describing hard nonlinearities, which are discontinuous nonlinearities that cannot be locally approximated by linear functions. One assumes the nonlinearity is driven by a sinusoid with a given frequency and magnitude, that the output can be described by an odd Fourier series (no DC component), and that the outputs are subsequently filtered by a linear portion of the system (in our case, the hand on the knob) with low-pass properties that eliminate all harmonics higher than the first harmonic so that only the first term in the series need be considered.

Figure 41 contains a block diagram of a system amenable to describing function analysis, with a reference input $r(t)$, plant input $u(t)$, system output, $y(t)$, nonlinearity output, $w(t)$, linear component, $G(j\omega)$, and nonlinear component, $H(j\omega)$. For basic describing function analysis, the reference input, $r(t)$, is normally assumed to be zero. The describing function is defined as the complex ratio of the fundamental component of

the nonlinear element to the input sinusoid. Describing functions may be obtained in three ways: analytically, by numerical integration, and experimentally. Slotine and Li (1991) provide a readable introduction to the topic.




Figure 41: Block diagram of a closed-loop linear system with linear component, $G(j\omega)$, and nonlinear component, $H(j\omega)$

Gelb and VanderVelde (1968) present an often-referenced, clear, and thorough coverage of the subject. Book-length lecture notes by Damen (2001) also provide good treatment of describing functions, with reference to techniques involving MATLAB. Mougenet and Hayward (1995) used describing function analysis to investigate limit cycle oscillations in an Animate Systems, Inc. force-controlled hydraulic valve/piston/sensor subsystem. Mougenet and Hayward examined four different nonlinearities, identified electromagnetic hysteresis in the servo valve as the source of limit cycles, and designed an analog (continuous) lead-lag compensator to quench the limit cycles.

7.2 Application of Describing Function Analysis to Haptic Barrier Limit Cycles

Describing function analysis has been developed for continuous systems. Unfortunately, haptic systems with computer controllers are inherently discrete systems, and the ZOH aspect of the discrete implementation is one of the central issues of concern,

preventing the direct application of describing function analysis. Though the method should theoretically be extensible to the discrete domain, the following analysis uses continuous-time describing function analysis, approximating the effects of the ZOH delay in continuous time by using the transport operator, $e^{j\omega t_0}$, with the phase shift $\theta = \omega t_0$, where t_0 is half the sample period.

Chapter 6 showed that for the cases of interest, limit cycle oscillations usually occur between the border of two encoder counts (see Figure 37 on Page 81). This situation, involving switching between just two output torque levels, offers the chance to use the describing function for a relay, a well-known nonlinearity in describing function analysis. The relay nonlinearity is a special case of the saturation nonlinearity with an infinitely small linear region. The relay output is either on (M) or off ($-M$). Figure 42 contains a graph of the relay nonlinearity.




Figure 42: The relay nonlinearity

One can easily obtain the describing function for the relay nonlinearity analytically (see Slotine and Li, p. 174):

$$N(A) = \frac{4M}{\pi A},$$

where $N(A)$ is the describing function, M is the magnitude of the relay nonlinearity, and A is the amplitude of the limit cycle.

The situation of a hand grasping a haptic knob and lightly pushing it against a virtual barrier deviates in two important ways from the assumptions of basic describing function analysis. First, since quantized displacement toggles between zero and a nonzero value, rather than between two values equal in magnitude and opposite in sign, the nonlinearity is actually more like a relay nonlinearity with a DC offset, violating the assumption that the DC component of the Fourier series can be neglected and that the nonlinearity is odd. Second, since the human applies a constant gentle bias torque to the knob, the reference input, $r(t)$, is nonzero. These two points cancel each other when one assumes that the user's bias torque is exactly

$$\mathbf{T}_{bias} = -\frac{1}{2}K\Delta\theta,$$

where K is the virtual barrier stiffness and $\Delta\theta$ is the displacement quantization.

To understand this, consider two cases, the first case where the knob is just short of barrier penetration, $\theta \geq 0$, and the encoder count equals 0, and the second case where the knob is just inside the barrier, $\theta < 0$, and the displacement reading equals $-\Delta\theta$. These two cases define a displacement quantization approximation good for all θ , which is a non-odd, nonlinearity, shown in Figure 43:




Figure 43: Virtual barrier nonlinearity

For negative values of θ , the quantized position signal will be $-\Delta\theta$, but the output of the nonlinearity for negative θ is shown in Figure 43 as $-K\Delta\theta$ as the result of an assumption for simplicity that the nonlinearity, $H(j\omega)$, also contains the stiffness constant, K , so that the output of the nonlinearity, $w(t)$, is in units of torque to match the reference input, $r(t)$, and the plant input, $u(t)$. Remember as well that the summation in the block diagram (Figure 41) reverses the sign of $w(t)$ so that the plant input, $u(t)$ is positive (torque pushing out of the barrier) for barrier penetrations indicated by negative θ .

If the person grasping the knob applies a bias torque of $T_{bias} = -\frac{1}{2}K\Delta\theta$, we have the following situation:

$$r(t) = -\frac{1}{2}K\Delta\theta$$

$$w(t) = \begin{cases} 0 & \theta \geq 0 \\ -K\Delta\theta & \theta < 0 \end{cases}$$

$$u(t) = \begin{cases} 0 & \theta \geq 0 \\ K\Delta\theta & \theta < 0 \end{cases}$$

which has a nonzero reference input, $r(t)$, and a non-odd nonlinearity.

If we take the reference input, assume it is constant, and include it in the nonlinear component so that its effect is apparent in the nonlinear output, $w'(t)$, we get:

$$r'(t) = 0$$

$$w'(t) = \begin{cases} \frac{1}{2}K\Delta\theta & \theta \geq 0 \\ -\frac{1}{2}K\Delta\theta & \theta < 0 \end{cases}$$

$$u'(t) = \begin{cases} 0 & \theta \geq 0 \\ \frac{1}{2}K\Delta\theta & \theta < 0 \end{cases}$$

where the prime notation on $r'(t)$, $w'(t)$, and $u'(t)$ indicates modification of the expressions, not differentiation. Accounting for the sign reversal in the summation, this amounts to adding $\frac{1}{2}K\Delta\theta$ to $w(t)$, yielding the following odd nonlinearity for $w'(t)$:




Figure 44: Virtual barrier nonlinearity with bias torque included

With the problem now set up in a fashion amenable to basic describing function analysis, we can turn our attention to the prediction of limit cycles. Through extension of the Nyquist stability criterion, one can predict that limit cycles will occur when

$$G(j\omega) = \frac{-1}{N(A)},$$

where $G(j\omega)$ is the open-loop transfer function of the linear element of the system. By plotting the negative inverse of the describing function along with the Nyquist plot of the linear element, one can obtain the predicted limit cycle oscillation magnitude and frequency simply by looking for the intersection of the two plots. Figure 45 contains a Nyquist plot of the linear system with the nominal describing function (dashed line) and delayed describing function (angled line with circle denoting intercept). A given sample rate is associated with a given ZOH delay, which will result in a particular counter-clockwise shift in the plot of the negative inverse describing function.




Figure 45: Nyquist plot of linear system with nominal describing function (dashed line) and delayed describing function (angled line with circle denoting intercept)

Once one has identified the point of intersection between the negative inverse describing function and the Nyquist plot, one knows the magnitude of the negative inverse describing function and can calculate the limit cycle magnitude. Starting with the relationship for the describing function:

$$N(A) = \frac{4M}{\pi A},$$

the negative inverse describing function is:

$$\frac{-1}{N(A)} = \frac{-\pi A}{4M},$$

and solving for the magnitude of A ,

$$|A| = \left| \frac{-1}{N(A)} \cdot \frac{-4M}{\pi} \right|$$

One can immediately see that the magnitude of the nonlinearity, M , is directly related to the magnitude of the limit cycle oscillations. For the haptic knob contacting a virtual barrier, $M = \frac{1}{2}K\Delta\theta$, and thus both the barrier stiffness and displacement quantization are directly related to limit cycle oscillation magnitude.

7.3 Results

Results were computed for the same seven sample rates, from 100 Hz to 10 kHz, as used in the simulations, giving seven different possible phase shifts for the negative inverse describing function plot. Figure 46 shows a portion of the Nyquist plot of the linear system along with seven negative inverse describing function plots, one for each delay. Larger angular shifts are associated with slower sample rates. One can see that for a given displacement resolution (constant M , given also a constant barrier stiffness of 0.27 Nm/rad), oscillation magnitude (related to the distance of the intercept point from the origin of the complex plane) increases for slower sample rates.




Figure 46: Nyquist plot (curve) with angled lines plotting describing functions for various ZOH delays, with circles indicating intersections

For each of the seven delays, results were computed for the same seven displacement resolutions (128 to 8192 counts/revolution) used in the simulations. Figure 47 shows the log magnitude of the limit cycles predicted by describing function analysis for the 49 cases in question. Figure 48 shows the same plot (left) next to the corresponding plot from the simulations (right) for ease of comparison. One can immediately see that describing function analysis does an excellent job of capturing the sensitivity of limit cycle oscillation magnitude to sample rate and displacement quantization; however, plotting the results on a log magnitude scale makes it difficult to examine differences in absolute magnitudes between the two techniques.




Figure 47: Log magnitude of limit cycles predicted by describing function analysis




Figure 48: Log magnitude of limit cycles predicted by describing function analysis (left) compared to Log magnitude of limit cycles predicted by simulations (right)

Figure 49 shows the difference between the DFA results and simulation magnitudes as a percentage of the simulation magnitudes. DFA consistently underestimates the limit cycle oscillation magnitude compared to the simulation results, despite the fact that both approaches consider the same linear system, same parameter ranges, and neglect other

nonlinearities such as friction. The differences plotted in Figure 49 range from -75 to -18% , averaging -55% , with a standard deviation of 15% .

The assumption used in describing function analysis that the linear system will filter out all but the first harmonic of the nonlinearity could potentially lead to an underestimation of oscillation magnitude if higher-order harmonics remain significant. To check this, the DFA was repeated with modifications to consider the first four harmonics of the describing function. Consideration of the additional harmonics changed the range of the differences between DFA and simulation to be between -71% and -4% and changed the mean and standard deviation to -47% and 18% , respectively. Adding the harmonics yielded magnitude estimates on average 47% under the simulation estimates, compared to 55% under for DFA with only the first harmonic considered; a significant difference, but not enough to explain the full amount of the difference observed between DFA and simulation.




Figure 49: Difference between DFA and simulation magnitudes as a percentage of simulation magnitudes

One other difference between the simulations and the describing function analysis is that the describing function analysis uses a transport delay, $e^{j\omega_0}$, which is only an approximation of the ZOH. The use of the transport operator, with a delay equal to half the sample period, assumes that every barrier threshold crossing (zero crossing) for barrier penetration and exit miraculously occurs at exactly halfway between sample instants. If barrier penetration times have a zero-mean stochastic variation around this midpoint, then the transport operator should closely approximate the ZOH effect; however if for some reason barrier threshold crossings happen on average more than halfway through a sample period, the use of the transport operator with half sample period delay would underestimate oscillation magnitude.

Figure 50 illustrates the frequencies of limit cycles predicted by DFA. Figure 51 shows the same plot (left) next to the frequency predictions from the simulations (right) for easy comparison. DFA and simulation provide nearly identical predictions of limit cycle frequency sensitivity – no significant sensitivity to displacement resolution and pronounced sensitivity to sample rate. The DFA plot also does not have the irregularities of the simulations; DFA was not subject to random numerical variations as the simulation predictions were.




Figure 50: Frequencies of limit cycles predicted by describing function analysis




Figure 51: Frequencies of limit cycles predicted by describing function analysis (left) compared to frequencies of limit cycles predicted by simulations (right)




Figure 52: Difference between DFA and simulation frequencies as a percentage of simulation frequencies

7.4 Discussion

Continuous describing function analysis using a relay nonlinearity with a transport delay does an excellent job of predicting the sensitivities of limit cycle magnitude and frequency to changes in sample rate and displacement resolution. DFA provides estimates of limit cycle oscillation magnitude well within a factor of two of the simulation predictions, consistently estimating smaller magnitudes than predicted with simulation. Consideration of higher-order harmonics provides a partial explanation of this underestimation of oscillation magnitude. DFA frequency predictions are quite close to simulation predictions (within 20%). One potential explanation for the fact that DFA frequency predictions were more accurate than magnitude predictions is that the Bode magnitude of the system response varies steeply with changes in frequency beyond the

natural frequency of the system; a fairly small change in frequency is associated with a large change in magnitude.

Given the simplicity and computational efficiency of limit cycle analysis, it provides an attractive tool compared to computationally expensive simulations which make repeated calls to differential equation solving routines. An entire cohort of 49 simulation cases took 1-3 hours to compute using Matlab software on a 400 MHz Pentium computer, whereas DFA of 49 cases took approximately 10 seconds. DFA offers a more rapid method than simulation to explore the sensitivity of limit cycle oscillation magnitude and frequency to variations in displacement quantization and sample rate. DFA could also be much more easily extended to consider different linear system components (e.g., stronger grip forces, physical damping, etc.).

Chapter 8 : Hardware Testing

8.1 Introduction

To validate the simulation results presented in the previous chapter, the HREF haptic knob system was used to simulate a virtual barrier for a range of four encoder resolutions (256 counts/revolution to 2048 counts/revolution) and four sample rates (455 Hz to 4 kHz). These two independent variables, with four levels each, created sixteen possible combinations of encoder resolution and sample rate, all of which were tested. This set of experimental cases is a subset of the 49 cases (two variables, seven levels each) used for the simulations.

8.2 Methods

Tests were conducted with one highly trained subject. The HREF haptic knob presented a unidirectional virtual barrier (a rotational “wall”) with a stiffness of 0.27 Nm/rad. Prior to each experimental trial, the zero-force voltage level output from the grip force load cell amplifier was noted for later subtraction from the measured grip force voltage. For each trial, the subject then grasped the knob using the simple pinch grasp studied in Chapter 4 (see Figure 2 on Page 25 for a photograph of this grasp). A strip-chart style display on the target computer displayed deviations from the target grip

force of 0.71 N and allowed the subject to achieve the target grasp force and maintain a constant grasp force once in contact with the virtual barrier. While watching the grip force display, the subject applied just enough torque to the knob to maintain contact with the virtual barrier. After 1-3 seconds of sustained buzzing, the subject released the knob and pressed a key on the host computer keyboard to stop the trial. The host computer then automatically uploaded the grasp force, commanded torque, and displacement data from the target computer.

8.3 Results

All but one of the 16 cases resulted in sustained oscillations. Appendix D contains snapshots of a portion of the raw data representing a period of sustained oscillations for each of these sixteen trials, with displacement and grip force plotted. The case with the highest encoder resolution (2048 counts/revolution) and the lowest sample rate (455 Hz) failed to exhibit sustained oscillations, displaying instead initial contact oscillations that rapidly died out.

Visual feedback of deviation from the target grip force allowed the subject to maintain the target grip force. Figure 53 illustrates mean grip force during the periods of oscillation under study for each of the 16 trials. The flatness of the surface demonstrates the subject's success at maintaining the target grip force. Table 8 contains actual grip force values, with a mean of 0.72 N and a standard deviation of 0.05 N (7%), compared to the target of 0.71 N. The raw data illustrated in Appendix D shows that the grip force remained quite constant during each interval studied.




Figure 53: Grip Force for Each of the Sixteen Trials

Table 8: Grip Forces for Each of the Sixteen Trials

		Sample Rate (Hz)				mean	std. dev.
		455	1000	2000	4000		
Encoder Resolution (bits/rev)	8	0.72	0.64	0.73	0.73	0.70	0.04
	9	0.73	0.79	0.76	0.76	0.76	0.02
	10	0.67	0.73	0.66	0.74	0.70	0.04
	11	0.69	0.64	0.74	0.78	0.71	0.06
mean		0.70	0.70	0.72	0.75	0.72	
std. dev.		0.03	0.07	0.04	0.02		0.05

Figure 54 shows thumbnail plots of raw displacement data for the oscillation volleys under study for the sixteen cases, with sample rate improving from left to right and encoder resolution improving from top to bottom. One can see that both worsening sample rate and worsening encoder resolution leads to larger limit cycle oscillations.

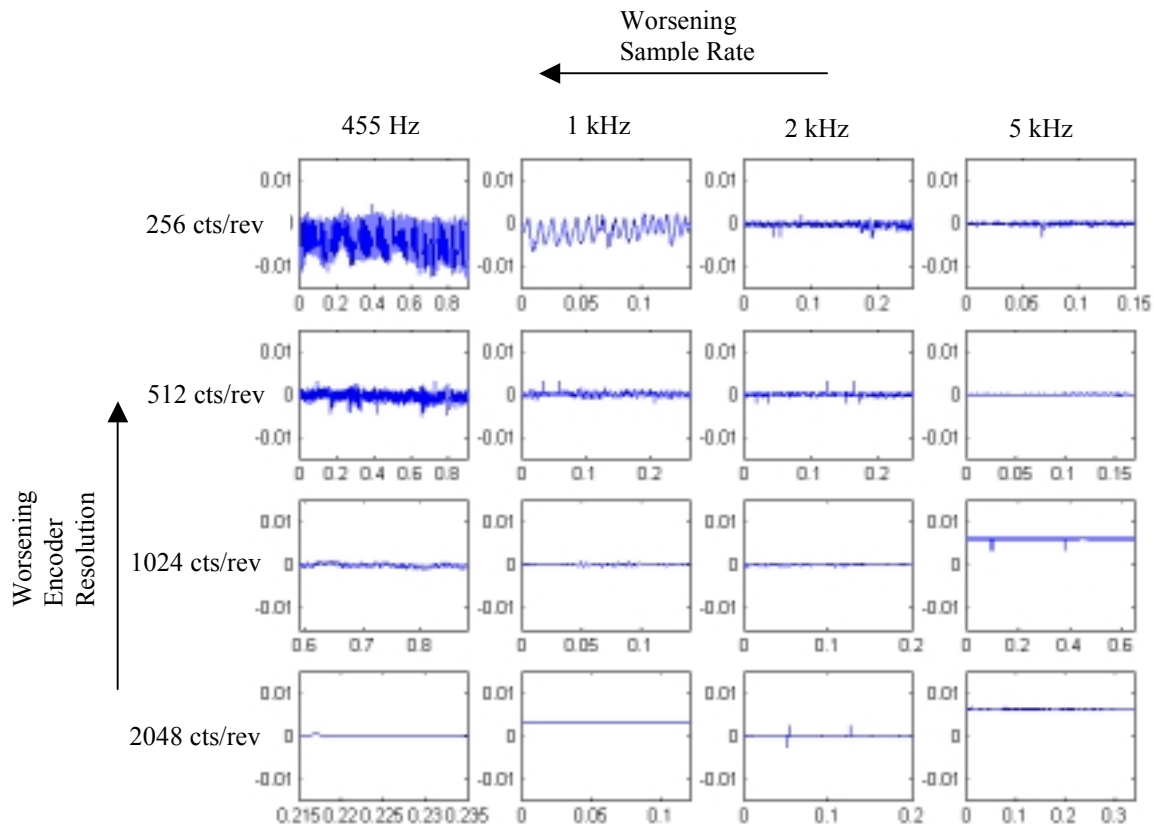


Figure 54: Thumbnail Sketches of Raw Displacement Data

Figure 55 shows the root-mean-square (RMS) magnitude of the limit cycles as a function of servo rate and encoder resolution. It shows very similar behavior to the simulation results represented in Figure 33 on Page 79. Figure 56 shows the same empirical data as Figure 55, but adjusted to represent peak-to-peak magnitude and plotted on a log magnitude scale. It compares well to the simulation results shown in Figure 34 on Page 79, and shows oscillation magnitude growing with both worsening sample rate and worsening encoder resolution.




Figure 55: RMS Magnitude of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution

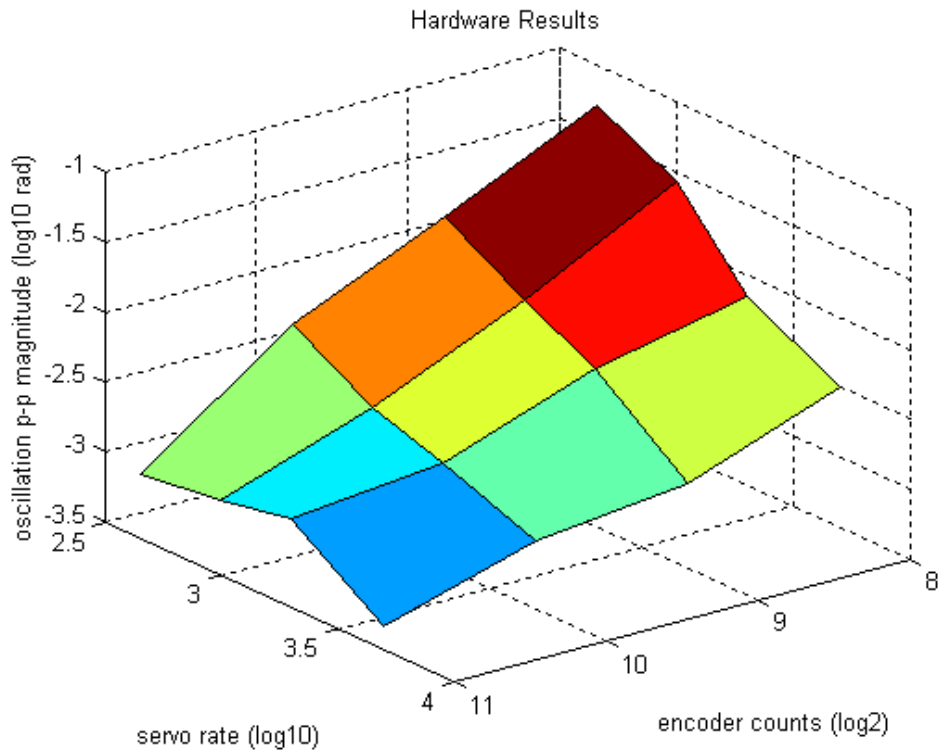


Figure 56: Log10 peak-to-peak Magnitude of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution

Figure 57 shows the hardware results (shaded mesh) superimposed along with the relevant simulation results (light mesh), emphasizing the excellent agreement of the two approaches with regards to sensitivity of oscillation magnitude to changes in sample rate and displacement resolution; however, the log magnitude plot acts to obscure absolute differences between the two sets of results. Figure 58 plots the difference between the hardware results and the simulation results, as a percentage of the simulation results. On average for the sixteen cases, hardware oscillations were 20% larger than simulation oscillations, with a range from 74% lower than the simulations to 150% larger than the simulations, and a standard deviation of 60%. By and large, oscillation magnitudes in the hardware test were within a factor of two of the simulation results.




Figure 57: A comparison of simulation to hardware results for oscillation magnitude as a function of sample rate and displacement resolution




Figure 58: Difference between hardware testing and simulation magnitudes as a percentage of simulation magnitudes

Figure 59 presents the principal oscillation frequency, determined by inspection of the data plotted in Appendix D. It reveals a trend towards increasing frequency with increasing sample rate, a trend also observed in the simulation results. The frequencies across all cases are approximately double those from the simulation results, and the empirical data shows a sensitivity to encoder resolution at low servo rates that the simulation results do not possess.




Figure 59: Principal Frequency of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution

8.4 Discussion

Methodical exploration of the sample rate and encoder resolution parameter space using haptic hardware presenting a virtual barrier to a human subject grasping a knob revealed that both sample rate and encoder resolution contribute to the magnitude of limit cycle oscillations upon contact with a virtual barrier. These empirical results agree well with simulations presented earlier. Improvements in sample rate or encoder resolution resulted in smaller oscillations.

Faster sample rates yielded oscillations at audibly higher frequencies. The oscillations for the higher resolution cases with the faster sample rates (particularly 2048 counts/revolution at 4 kHz) were sometimes more audible than perceivable tactually. The oscillation frequencies measured in hardware testing differed significantly

from simulations, with a tendency towards higher frequencies overall and a sensitivity to displacement resolution at high sample rates that was not apparent in the simulations.

The most likely explanation for this deviation is the fact that bias torque was not controlled for the hardware experiments, and in simulation, the oscillation frequency was noted to vary considerably with changes in bias torque.

Chapter 9 : Effects of Sample Rate and Displacement Quantization on the Amplitude and Frequency of Limit Cycle Oscillations

Work in earlier chapters has shown that a displacement quantization cannot by itself add energy to a haptic interface with a unilateral virtual barrier, but that when present in a sampled data system with zero-order hold (ZOH) torque output, displacement quantization acts to worsen the energy-instilling effects of the ZOH identified by earlier investigators. This chapter will explain how this interaction arises, and how the two system design parameters of sample rate and displacement quantization affect the system quality metrics: limit cycle magnitude and oscillation frequency.

9.1 Why Does Displacement Quantization Worsen ZOH-Induced Energy Leaks?

For a virtual environment defined with positive displacement pointing away from a virtual barrier and negative displacement pointing into it, encoder quantization acts like a floor function, rounding the measured displacement down to the nearest multiple of encoder resolution (e.g., from -0.5 down to -1 , if $\Delta\theta = 1$). For a manipulandum moving anywhere within the region bounded by a “zero” encoder reading, no barrier torque will result. Once the manipulandum just crosses the lower threshold of this range, a “-1”

encoder count will be recognized, and the controller will command a torque $T=K\Delta\theta$, where K is the spring constant, and $\Delta\theta$ is the encoder quantum. When the manipulandum passes just below the bottom edge of the “-1” encoder count region to register a “-2,” the barrier torque immediately jumps to $T=2K\Delta\theta$, and so on. The round-down (floor) function means that the quantized barrier penetration will *always* be greater than or equal to the actual penetration.

By always making the torque greater with quantization than without (both during penetration and during withdrawal), displacement quantization effectively increases spring torque by an offset on average equal to half the encoder quantum times the spring stiffness. Since ZOH-induced energy leak magnitudes are a function of spring torque, the displacement quantization acts to worsen energy leaks for systems with ZOH. In the case of a continuous-time system with displacement quantization, there would be no ZOH leak to worsen by increasing the barrier stiffness. The inverse case is different; as previous investigators have shown, energy leaks can exist for systems with ZOH in the absence of displacement quantization. For systems with both ZOH and displacement quantization, the leak is proportional both to the stiffness of the barrier spring and the magnitude of the delay induced by the ZOH.

Figure 37 on Page 81 demonstrated that for the great majority of the cases studied, limit cycle oscillations remained smaller than +/- 1 encoder count. This permits a simplified analysis, where the commanded torque resulting from the quantized encoder position input to a unilateral spring can only have two values: zero and $K\Delta\theta$, where $\Delta\theta$ is a constant denoting the encoder step size (quantization interval).

Figure 60 contains an illustration of a case where the knob penetrates the barrier by less than one encoder quantum and then exits the barrier five time steps later. The curve represents the continuous time displacement trajectory. The dotted line represents the commanded torque output for a system with ZOH and infinite resolution displacement sensing (scaled similarly to previous plots). The dashed line represents the commanded torque output for a system with ZOH and quantized displacement sensing. Time interval Δt_a denotes the time between actual barrier penetration and the update of the barrier torque (to $K\Delta\theta$) by the controller. Time interval Δt_b denotes the time between actual barrier exit and the update of the barrier torque (to zero) by the controller.




Figure 60: Illustration of Barrier Penetration and Resultant Torque Outputs for a Traditional ZOH System and a ZOH System with Displacement Quantization

During the interval, t_a , between barrier penetration and output torque updating, both discrete controllers behave the same way, keeping the torque at zero and allowing “free” penetration of the barrier with no resultant output torque. For controller updates inside the barrier, the continuous displacement controller exhibits the familiar lagging ZOH staircase. Since the penetration never exceeds one encoder count, the quantized output remains at the same (stronger) level for the duration of the penetration. In the “Symmetric Interval,” the quantized controller exerts the same torque during penetration as during withdrawal, adding no net energy. During the “Exit Period” the quantized controller exerts considerably more torque than the continuous displacement controller.

While the quantized controller adds more energy during the Exit Period, during the “Symmetric Interval” when the quantized controller is adding no energy, the continuous-displacement ZOH controller adds energy represented by the lightly shaded staircase areas. Which controller adds more net energy to the system? One can inspect the error torque-time integral difference between the cases by comparing the areas of the lightly shaded staircase triangles to the heavily shaded rectangle. The triangles occupy 38 grid squares, while the rectangle occupies only 30 grid squares. Since we know that the quantized controller actually instills more energy in the system, erroneous torque integrated over time cannot be an accurate measure of energy instilled.

Mechanical work is performed by applying a force (or torque) over a displacement, so energy is added to the system according to the energy error integral equation:

$$E_{leak} = \int_{t_1}^{t_2} T_{error}(t) \dot{\theta}(t) dt ,$$

where $\dot{\theta}(t)$ is the angular velocity of the knob. We can approximate the energy gains for the two cases by inspection, using the formula

$$E_{leak} = \sum_{k=1}^N T_{error}(k) \dot{\theta}(k),$$

where $d\theta_k$ indicates the distance travelled during a given period. For the continuous displacement ZOH controller's four triangular areas add the following energy (estimating torque magnitude for the triangles at $\frac{1}{2}$ the maximum error):

$$2 \cdot 4 + .75 \cdot 1.5 + .75 \cdot 1.5 + 2 \cdot 4 = 18.5 \text{ energy units}$$

where each product term multiplies $\frac{1}{2}$ the maximum erroneous torque in units on the vertical axis (which represents torque for the command trajectories) times the distance travelled by the knob during the interval, in units along the vertical axis (which represents displacement for the knob trajectory). The quantized ZOH controller's extra push during the Exit Period is defined by the product of the extra torque, approximately 7.5 extra units on the vertical axis times the distance travelled by the knob during the period, approximately 5.5 units on the vertical axis, for an energy contribution of:

$$7.5 \cdot 5.5 = 41.25 \text{ energy units}$$

The quantized ZOH controller adds its erroneous torque pulse at precisely the worst moment, in the same direction as the exit velocity vector, with the same sign, effectively maximizing the work performed by the pulse on the knob!

9.2 Why Do the Bouncing Ball Simulations Show Less Sensitivity to Encoder Resolution Than the Hand-on-Knob Simulations?

If the simulation for a hand on a manipulandum reveals strong sensitivity for the magnitude of limit cycle oscillations with variations in both sample rate and encoder resolution, why does the simulation for the bouncing ball case show no sensitivity to encoder resolution at low sample rates and only modest sensitivity at high sample rates? (Figure 61, below, reprints the simulation plots for both cases for easy comparison.)

Since the contribution of encoder quantization to energy growth occurs at the threshold crossings, one would expect sensitivity to track with the frequency of threshold crossings. For sample rates less than a few thousand Hertz, the quadratic growth rates for the bouncing ball simulations were very high, resulting in increasingly deep penetrations of the wall followed by long ballistic trajectories above the wall, with infrequent threshold crossings. During deep penetrations across multiple encoder intervals, the quantized displacement ZOH controller's finite addition to the spring torque and its energy contribution during the exit period become less significant than the ZOH staircase errors that are present even without displacement quantization.

As a result of the insignificance of the threshold crossings to the ZOH leaks picked up on the deep wall penetrations, the bouncing ball simulations have undetectable sensitivity to encoder quantization. For higher servo rates where the ZOH leaks are less, and slower quadratic growth means more frequent bounces (thus more threshold crossings), the quadratic growth rate becomes modestly sensitive to encoder quantization.




Figure 61: Quadratic Energy Growth Rates for the Bouncing Ball Case (left) and Peak-to-Peak Magnitudes for the Hand-on-Knob Case (right)

9.3 Prediction of Limit Cycle Oscillation Magnitude

Earlier sections of this dissertation have shown that displacement quantization acts to multiply the effects of zero-order hold in producing limit cycle instilling energy leaks.

Actual limit cycle magnitude depends on the response of a second-order system subjected to unilateral control law changes, with work applied in impulses upon threshold crossing.

These work impulses depend on the magnitude of the errant torques and the distance traveled during the short time interval when torques are applied (which implies a reliance on velocity).

As shown previously, an estimate of the errant torque-time interval would not strictly estimate the energy leakage, since increasing torques affects the velocity, and

one really wants to know the integral $E_{leak} = \int_{t_1}^{t_2} T_{error}(t)\dot{\theta}(t)dt$. But what if the errant

torque-time interval, which is related to the product of displacement quantization and

sample period, could be used to form a first-order approximation of limit cycle

magnitude? This would satisfy observations that displacement quantization and sample

interval tend to have a multiplicative effect on limit cycle magnitude.

Defining limit cycle oscillations as:

$$\theta(t) = A \sin(\omega t),$$

the approximation is stated:

$$A = C \cdot \Delta\theta \cdot \Delta t,$$

where A is oscillation magnitude, C is a constant, $\Delta\theta$ is the displacement quantum and Δt is the sample period. Using a constant value of $C=500$, the limit cycle oscillation magnitude was estimated for the range of displacement quantizations and sample rates used for the simulation study, and then for the range used for the empirical testing with the HREF hardware. The approximation is remarkably accurate.

Figure 62 shows the simulation results for oscillation magnitude repeated from Figure 35 on Page 80. Figure 63 shows the output of the prediction approximation described above. The two plots are nearly identical, strongly supporting the assertion that the prediction approximation can estimate limit cycle magnitude. Figure 64 repeats the hardware results from Figure 56 on Page 103, and Figure 65 shows the output of the prediction approximation, also using the value $C=500$, but applying a $.707/2$ correction to the approximation to convert from peak-to-peak figures to RMS figures for comparison with the hardware results. The prediction agrees quite well with the hardware results, though not as well as with the simulation results. The prediction consistently overestimates limit cycle oscillations, perhaps because it neglects dissipative nonlinearities in the hardware such as Coulomb friction.




Figure 62: Simulation Results for Peak-to-Peak Oscillation Magnitude




Figure 63: Predictions for Peak-to-Peak Oscillation Magnitude




Figure 64: Log10 RMS Magnitude of Limit Cycle Oscillations as a Function of Servo Rate and Encoder Resolution




Figure 65: Log10 RMS Magnitude of Predictor Function Over Same Range of Parameters as Hardware Trials

9.4 Frequency Behavior

Simulation results from Chapter 6 showed that limit cycle oscillation frequency is quite sensitive to increases in sample rate, but relatively insensitive to increases in displacement resolution. Figure 66 illustrates the limit cycle frequency as a function of sample rate, averaged across cases with different displacement resolutions.




Figure 66: Average Limit Cycle Frequency as a Function of Servo Rate

The high sensitivity of limit cycle frequency to sample rate occurs because shortening the sample interval shortens the average time between actual barrier crossings and updating of the control law, shortening the "free" penetration interval and the erroneous torque pulse upon barrier exit. Shorter pulses instill less energy and accelerate the knob for a shorter time, allowing it to begin decelerating and to return to the zero crossing sooner. The duration of the erroneous torque signal is much smaller than the time constant of the hand/knob second-order system, which does not have time to reach steady state. One could view the end of the torque pulse as the interruption of a step response, allowing the system to return to zero displacement much sooner. The shorter the pulse, the quicker the return to zero displacement, and vice versa, so short pulses

(caused by faster sample rates) correspond to higher frequency limit cycles, and longer pulses (caused by slower sample rates) correspond to lower frequency limit cycles.

9.5 The Role of Friction

One of the many ways in which the empirical testing with the HREF hardware differs from the simulations is that the hardware presents friction, which was not modeled in the simulations. Dynamometer testing of the HREF hardware revealed Coulomb friction (a negative constant times the sign of the velocity) of 0.00056 N-m (0.56 mN-m). This friction can have a beneficial effect by quenching limit cycles when oscillations of the commanded torque are small. This effect may help to explain why it was hard to excite oscillations for very high encoder resolutions during the hardware tests.

Niemeyer (2001) has suggested specifying the displacement quantization and maximum virtual spring stiffness so that the quantum change in torque due to a change of one encoder count is below the Coulomb friction level. With parameters thus set, oscillation between two adjacent encoder counts will produce no torque output, and these oscillations will not be sustained. For the HREF system with Coulomb friction ($T_{Coulomb}$) of 0.00056 N-m and a maximum stiffness (K) of 0.27 N-m/rad, the required resolution ($\Delta\theta$) would be:

$$\Delta\theta = \frac{T_{Coulomb}}{K} = \frac{0.00056 \text{ Nm}}{0.27 \text{ Nm/rad}} = 2.07 \times 10^{-3} \text{ rad}$$

which equates to 3,030 counts/revolution, a reasonable design goal in many instances.

Limit cycle oscillations at resolutions better than this were, in fact, quite hard to excite with the HREF hardware.

Chapter 10 : Design Implications

10.1 Introduction

The primary goal of this work is to understand the effect of displacement resolution and sample rate on limit cycle oscillation magnitude and frequency. An understanding of this relationship can contribute both to the design of haptic systems with acceptable limit cycle behavior and to the implementation of methods to mitigate limit cycles. Previous chapters developed validated dynamic models and explanations for limit cycle behavior, showing that increasing sample rate acts to both decrease limit cycle oscillation magnitude while increasing oscillation frequency, and that increasing displacement resolution primarily acts to decrease oscillation magnitude without much effect on frequency. The knowledge presented in previous chapters becomes useful when one considers its implications for the design of haptic systems. This chapter examines those design implications.

After enumerating the design variables (including displacement resolution and sample rate) that can impact limit cycle behavior, this chapter will explore the role of threshold and “loudness” perception for the senses of touch and hearing, which result in a design bias towards decreasing limit cycle oscillation magnitude without increasing

oscillation frequency. The chapter will then discuss some methods suggested by earlier authors for mitigating ZOH-induced limit cycle oscillations, will present an example of the inadequacy of increasing sampling rate to abate limit cycles, will cover three more liabilities of excessively fast sampling rates, and will conclude with a suggested design approach.

10.2 Design Variables Affecting Limit Cycle Behavior

Many design variables impact limit cycle oscillation magnitude and frequency. The pursuit of a haptic system with satisfactory limit cycle behavior is a classic systems engineering problem. Changes in any design variable to improve limit cycle behavior will come with a cost that could be financial or could impact another system performance metric such as wall stiffness. Some key variables affecting limit cycle behavior appear below, followed by the next section, which discusses human thresholds for touch and hearing perception and the bearing that human perception has on the design problem.

1. Displacement resolution – As shown in previous chapters, improving displacement resolution reduces the magnitude of limit cycle oscillations without raising oscillation frequency. Improving resolution generally increases system cost and can sometimes increase space and power requirements.
2. Sample rate – As shown in previous chapters, increasing sample rate decreases limit cycle oscillation magnitude but increases limit cycle oscillation frequency, which can make limit cycles more easily perceptible if it raises the oscillations to a frequency where the human sense of touch or hearing is more sensitive.

3. Virtual barrier stiffness – This dissertation and the work of many others have shown that limit cycle oscillation magnitude scales directly with barrier (wall) stiffness. One can prevent the appearance of limit cycles by limiting barrier stiffness to some maximum value; however, this limitation on barrier stiffness restricts the capability of the haptic device to present convincing barriers to the user. Conversely, improvements in other variables such as displacement resolution and sample rate allow haptic systems to have stiffer barriers without exciting unwanted limit cycles.
4. Physical damping -- Physical damping can dissipate energy and reduce limit cycle magnitude while reducing the damped natural frequency of a system; however, physical damping can be difficult to add repeatably to manufactured haptic systems, and can limit the systems' ability to present delicate haptic sensations. Physical damping can be reversed with negative virtual damping, but this can limit the maximum allowable wall stiffness (Colgate and Schenkel, 1994).
5. Physical inertia – The addition of physical inertia will lower the natural frequency of a system and reduce the magnitude of oscillations; however, in a system with limited actuator authority, high inertia can make it difficult to present high frequency haptic sensations to the user. Though not the most prominent issue in direct-drive haptic knobs, if inertia is too high, the user may also find the inertia distasteful. If actuator authority, bandwidth requirements, and user preferences permit, the addition of some inertia may be helpful.

6. Velocity signal quality – If there is no separate velocity sensor, this variable tracks with displacement resolution, as the most common way of obtaining velocity signals for a haptic device is to differentiate the displacement signal. The quality of a differentiated velocity signal degrades quickly with worsening displacement resolution. If independent velocity information can be obtained through a separate sensor, this velocity signal can be used for virtual damping to decrease limit cycle oscillation magnitude, though as mentioned above, the addition of virtual damping limits the maximum allowable wall stiffness (Colgate and Schenkel, 1994).
7. Grasp impedance – As shown in the system identification experiments presented in this dissertation, increasing grasp force will result in increasing grasp impedance (damping and stiffness). Though not a design variable per se, the designer should be aware that light grasps, with minimal damping and stiffness, have the potential to excite the highest magnitude limit cycle oscillations.

10.3 Touch and Hearing

Haptics researchers know well that the sense of touch displays a sharp peak in sensitivity (a threshold trough) at 250 Hz (Verrillo, 1968). Between 10 and 250 Hz, the magnitude of displacement oscillations required to generate a perceptible sensation drops rapidly. The tactile sensation threshold begins rising above 250 Hz as humans become less sensitive to vibrations at higher frequencies. The inclination to design for limit cycle oscillations well below or well above 250 Hz is quickly abated when one considers that as frequency rises above 250 Hz and tactile sensitivity decreases, aural sensitivity rapidly *increases*. Recall that for some of the hardware testing done in Chapter 8, at some of the

higher sample rates and higher encoder resolutions, the limit cycle oscillations were more audible than tactually perceptible.

Thresholds and perceived loudness levels for the sense of hearing have been extensively studied (Katz, 1985; ISO, 1961; Scharf and Buus, 1986). The sound pressure levels at which humans perceive sounds of equal magnitude can be defined as a family of “equal loudness contours,” with each contour being a curve defining a perceived loudness level, starting from a 1 kHz tone of a given amplitude. Fletcher and Munson first defined equal loudness contours in the 1930s at Bell Labs. Figure 67 contains the classic Fletcher-Munson equal loudness contours. The lowest member of the family of curves corresponds to hearing threshold (with 0 dB defined as the sound pressure level at which one can just hear a 1 kHz pure tone).

Equal loudness contours show a pronounced notch at 3 kHz, where the human ear is most sensitive. The ear is slightly less sensitive at 1.5 kHz and 5 kHz. Sensitivity remains constant from about 1.5 kHz down to about 600 Hz. Sensitivity drops off sharply below 600 Hz and above 5 kHz. The range of oscillation frequencies between 250 Hz and 3 kHz seems quite unattractive, since it is bounded by the peak sensitivity regions of the tactile and aural senses, respectively. The ear also remains sensitive in this range, and typical haptic devices are quite efficient at generating audible vibrations in this range if properly excited. The possibility of increasing limit cycle oscillations with coarse encoders and fast sampling to a frequency high enough that it could not be heard is tenuous. By process of elimination, that leaves the range below 100 Hz as the least offensive frequency range for limit cycle oscillation.




Figure 67: Equal Loudness Contours for the Sense of Hearing, Based On the Work of Fletcher and Munson (image: D. Cross, in Gotfrit 2001, by permission)

10.4 Design Implications

One of the notable features about haptic wall contact with low resolution encoders is that meaningful velocity information is not available. Limit cycle oscillations very often occur only between one encoder count and the next. In other situations, the oscillations only cross several encoder counts. In these cases neither velocity determination by counting pulses in a period of time nor counting the time interval between encoder pulses can provide useful velocity information. This absence of velocity information renders several previously proposed approaches useless.

Gillespie proposed two approaches for reducing the effect of the zero-order hold (Gillespie and Cutkosky, 1996). The first, a model-based half-sample prediction controller, requires accurate knowledge of the model state (position and velocity) in order to simulate ahead half a time step. For a typical system implementation with an optical

encoder, displacement quantization reduces the quality of the position signal and prevents derivation of a meaningful velocity signal, eliminating a model-based half-sample prediction controller as a viable approach. Gillespie's second proposed approach, a controller designed in the digital domain, similarly relies on full state feedback and cannot work without reasonable position and velocity signals.

Ellis et al. suggested using predictive force output calculations based upon the difference between the last force output and the current force output to drive a zero-order hold system to approximate the force-time integral of a first-order hold system (e.g., in the case of a rising force, the controller sets the ZOH force output a bit higher than it otherwise would, so that during the sample interval, the area under the force-time curve is the same as it would have been for first-order hold). Ellis' numerical simulations do not appear to consider displacement quantization, and his empirical confirmation uses a Hayward Pantograph (Hayward et al., 1994) with potentiometers sensing joint displacement (though presumably Ellis' system was still subject to 12 or 16-bit quantization in the A/D conversion). The extreme quantization in force (torque) output for systems with coarse displacement sensing would render useless techniques such as first-order hold that depend on predictions of the rate of change of the output force (torque).

10.5 An Example to Discourage High Sample Rates and Low Displacement Resolutions.

The BMW iDrive system uses a 7.4:1 transmission to allow the use of a 360-count (after quadrature) optical encoder to create 2650 counts per revolution of the knob. Sensor transmissions such as this occupy precious space and cost money. A direct-drive

sensor would allow more design flexibility, smaller packages, and lower cost. Results presented previously have shown that increasing sample rate can decrease limit cycle oscillation magnitude, but is likely to increase oscillation frequency. Figure 68 illustrates a simulation with two cases: the current iDrive parameters of 2650 counts/rev and 1 kHz sample rate, and a direct-drive sensor test case with 360 counts/rev and 20 kHz sample rate that promises to decrease limit cycle oscillation magnitude by using a very high sample rate, despite the coarse encoder.




Figure 68: Simulation Results Showing Oscillation Magnitude for Current Design (left corner, 2650 count/rev, 1 kHz) and Test Case (right corner, 360 counts/rev, 20 kHz)

A virtual barrier was implemented on the HREF knob system with a sample rate of 25 kHz and an encoder resolution of 360 counts/revolution. Oscillations for this test case did in fact seem below tactually perceptible limits, but the system **whined** loudly.

Beyond the effects of limit cycle oscillations a second noxious percept was noted: perceptible steps in spring force as the knob crossed encoder counts on the way into or out of the barrier. These perceptible steps also existed for some higher resolutions, suggesting that limit cycle oscillations are not the sole perceptual driver for specification of displacement resolution.

Figure 69 shows limit cycle oscillations recorded during this test, with a peak-to-peak magnitude of 0.0014 radians and a frequency of 2200 Hz (approximately one-tenth the sample rate). The square signal represents the output torque (scaled by the spring constant to indicate encoder quantization, as has been done in previous plots) and the curve represents actual encoder position, as measured by the 640,000 counts/revolution encoder, with the dots embedded along the curve representing sampling instants. One can see the classic effects of the control law asynchrony identified by Gillespie, with control law updates occurring coincident with the sample instants, but some time after the knob has penetrated or exited the barrier. Note that the oscillations are small enough to occur right on the threshold between two encoder counts – absolutely no velocity information is available, and the resulting quantization in the commanded torque output limits the usefulness of first-order hold techniques.




Figure 69: Limit Cycle Oscillations for Virtual Barrier at 25 kHz Sample Rate with 360 counts/revolution Encoder, Recorded from the HREF Hardware.

10.6 Liabilities of Sampling Too Quickly

Due to the minimum servo rate requirements typical of haptic systems, and the fact that haptics designers very often find themselves struggling with interfaces to virtual environments with relatively slow update rates, or with systems with limited processing resources, it is easy to adopt the sentiment that “faster is always better” when it comes to haptics servo rates. Perhaps counterintuitively, this is not always the case. Beyond the reasons already expressed in this work that discourage fast sample rates, there are at least three other examples of situations where fast sampling can cause trouble (Salcudean, 2001).

1. Finite Encoder Resolution – With quantized position sensors such as optical encoders, and velocity determined by $\Delta x / \Delta t$, the encoder value (x) may remain constant for many time samples and then “jump” to a new quantized value, creating a sudden increase in velocity (and an even worse artifact in acceleration). For a smaller Δt , this artifact will be

more severe. If recognized, appropriate low-pass filtering of the encoder signal before differentiation can mitigate this problem.

2. Dynamics Discretization – Discretization of a continuous system at high sample rates can introduce non-minimum phase zeros (discrete equivalents to zeros in the right-half s plane) that make it difficult to invert and control the plant. A compensator that is intended to be the complement of the plant will have the unstable poles, which would make it troublesome to deal with. Ideally, the complementary compensator would cancel the zeros in the plant, but the reality would fall short of the ideal.

3. Numerical Sensitivity -- For very high sample rates, filter and controller coefficients may be very small, which can make the system subject to numerical quantization problems.

10.7 A Suggested Approach

Though increasing sample rate can reduce limit cycle oscillation magnitude, it can also increase oscillations into easily audible frequency ranges. Since increasing displacement resolution decreases oscillation magnitude without substantially affecting oscillation frequency, displacement resolution should be the primary method for minimizing limit cycle oscillation magnitude. One should select the maximum sample rate that avoids easily audible oscillation frequencies, and then select the displacement resolution to bring oscillation magnitude below detectable levels. In some cases, it may be worthwhile to consider *lowering* very fast sample rates if their increased audibility outweighs any reductions they bring about in oscillation magnitude.

Figure 70 shows a notional optimization surface constructed by normalizing the logarithmic oscillation magnitude mesh and the oscillation frequency mesh from the simulations and creating a single mesh using a “maximum” function. In this mesh, smaller values represent the most attractive outcome (avoiding high oscillation magnitudes and high oscillation frequencies). This optimization surface is notional, and represents arbitrary weightings for oscillation magnitude and frequency. Yet it illustrates the suggested design approach.

Oscillation magnitude dominates the left portion of the mesh in Figure 70, and oscillation frequency dominates the right portion of the mesh. A designer would want to ride in the trough, avoiding low sample rates that would cause large magnitude oscillations, and avoiding high sample rates that would cause high frequency oscillations, and then using improvements in displacement resolution to gain additional advantage. The flat portion of the mesh in the foreground represents the regime where encoder resolution is high enough that the quantum change in torque due to a change of one encoder count is below the Coulomb friction level (as explained in Section 9.5), so increasing resolution further brings no benefits. Since the range of sample rates chosen for investigation extend above and below common the common sample rate for haptic systems of 1 kHz, the trough illustrated here runs right through 1 kHz (10^3 Hz).




Figure 70: Notional Optimization Surface

The above plot does not consider specific sensory thresholds or loudness contours for touch or hearing. A sophisticated design guide might offer a sensory-weighted optimization function to better predict the perceived noxiousness of chatter caused by ZOH and displacement quantization. The “complete” sensory-weighted optimization function would include weightings for every step in the chain from the inputs (displacement resolution and sample rate) to the output (perceived chatter noxiousness). It would consider not just the dynamic response of the hand/knob/actuator system, but also the audible frequency response of the system, and overlapping sensory weightings for the sense of touch and the sense of hearing. The audible frequency response of the system would need to be considered separately because a host of factors such as the sound-generating surfaces, sound paths, and device enclosure could make the frequency profile of the sound energy reaching the ears different from the profile reaching the

fingertips. Though interesting, such a sensory-weighted optimization function would be laborious to generate, would pertain only to a specific hardware implementation; the costs would likely outweigh any benefits obtained over simpler methods such as picking the maximum inaudible sample rate and then working to increase displacement resolution.

10.8 Other Design Implications

The suggested approach of picking an optimal sample rate and then seeking to increase displacement resolution until limit cycles become imperceptible largely reflects the current practice of commercial haptic system designers. Due to the relatively high cost of moderate to high resolution displacement sensors, relative to the extreme cost pressures of mass market commercial design, designers continue to feel constraints that are manifested in lower attainable barrier stiffnesses or additional complexity such as transmissions to increase the resolution available from inexpensive coarse sensors. These pressures, and the demonstrated importance of adequate displacement resolution, suggest that for high-volume commercial haptics applications that demand high fidelity, it may be worth a significant investment to develop low-cost, high-resolution displacement sensors.

Chapter 11 : Limit Cycle Mitigation Approaches

Earlier chapters developed validated dynamic models of the human hand grasping a haptic knob and presented explanations of the limit cycle behavior of haptic knobs in contact with virtual barriers. The previous chapter explored the design implications of this knowledge. This chapter examines how to make further use of the knowledge by informing potential approaches to mitigating limit cycle oscillations. Previous chapters have shown that low displacement resolution can worsen the magnitude of ZOH-induced limit cycle oscillations, and that efforts to reduce the magnitude of these oscillations must avoid increases in oscillation frequency that would make the limit cycles audible. The strong need for higher resolution displacement sensing and the desire to avoid near-term costs of better off-the-shelf sensors or long-term costs of developing inexpensive high-resolution sensors motivates attempts to devise control strategies that mitigate limit cycle oscillations without requiring increased displacement resolution. The absence of state information (indeed, the utter absence of velocity information) makes this task difficult.

There exist many candidate approaches for mitigating limit cycle oscillations, among them:

Increasing displacement resolution: This is the “obvious” solution, and would result in reduced limit cycle oscillation magnitude without increasing oscillation frequency. It would also enable higher-quality velocity signals for velocity-dependent effects such as virtual damping.

Adding physical damping or Coulomb friction: The addition of physical damping to a haptic mechanism (e.g., using viscous greases) would dissipate energy and reduce limit cycle oscillation magnitude. While physical damping and Coulomb friction can be attractive to a designer, in practice they can be difficult to add to a manufactured haptic device in a consistent manner. Coulomb friction greater than the product of the encoder quantum and maximum wall stiffness can effectively eliminate limit cycle oscillations. Both physical damping and Coulomb friction have the disadvantage that they make it difficult to feel delicate haptic sensations (lowering the dynamic range, in the case of Coulomb friction). Virtual negative damping can be added to counteract physical damping to allow more delicate sensations to be presented, but Colgate and Schenkel (1994) developed passivity conditions that show that the addition of virtual damping reduces the maximum barrier stiffness that can be achieved while remaining passive (see description in the beginning of Chapter 6).

Electromechanical damping: Techniques such as shorting motor windings together or using eddy current damping may provide opportunities for damping specific haptic systems. This effect would be different than adding physical damping because it could be disabled electronically, without requiring negative virtual damping, and it would be easier to produce repeatably.

Virtual damping using a velocity sensor: A haptic system with coarse displacement sensing and a separate velocity sensor could take advantage of the velocity sensor to provide virtual damping without relying on differentiation of the displacement signal. This approach would have to compete with the option of spending more money to simply increase

displacement sensor resolution. The addition of a second sensor adds cost, packaging and space constraints, and data acquisition complexity (additional digital inputs or analog-to-digital conversion would be required).

Corrective torque pulses: This technique, described in this chapter as “Impulse Damping” would apply torque pulses designed to compensate for torque errors upon entry and exit from a virtual barrier. A preliminary attempt to use this technique, presented later in the chapter, did not give good results, but further refinements may be worth pursuing.

Phase estimation damping: This technique, also described in more detail in this chapter, monitors zero crossings and attempts to present a torque wave out of phase with an estimated velocity sinusoid in order to generate an estimated damping command. A preliminary attempt to use this technique, presented later in the chapter, did not give good results, but further refinements may be worth pursuing.

Nonlinear virtual barriers: Nonlinear virtual walls or barriers with stiffness coefficients that increase with greater wall penetration require stronger grip forces to maintain barrier penetration, so as the knob penetrates farther into the barrier, the user's grip must become stronger, and thus the impedance of the user's grasp will become higher as the barrier stiffness increases. As wall penetration deepens, the effect of increasing grasp impedance acts to counter the effect of increasing wall stiffness, reducing susceptibility to limit cycles.

Two methods are discussed in this chapter, both of which achieved only modest success, but which may be worth further investigation. The first method, Impulse Damping, provides short force or torque impulses shortly after the device crosses from one encoder count to the next. These impulses are designed to cancel force/torque errors resulting from the combination of encoder quantization and zero-order hold effects. Several alternate approaches to this method are discussed. The second method, Phase

Estimation Damping, uses zero-crossing detection techniques to estimate the phase of limit cycle oscillations and generate a damping waveform 180 degrees out of phase with the estimated velocity of the limit cycle oscillations.

11.1 Typical Oscillations

Before describing the two mitigation schemes, let us examine the time response and phase plane plot for a typical limit cycle case at 2650 counts/revolution and 1 kHz sampling rate. Figure 71 shows the time response and Figure 72 shows the phase plane response, with manipulandum displacement on the horizontal axis and velocity on the vertical axis. The dotted vertical line in the middle of the phase plane plot denotes the position of the barrier threshold. Note the sharp inflection points that indicate changes between the inside barrier control law and the outside barrier control law, and that these inflection points do not occur precisely at the barrier. They have a phase delay, and correspond to the first sampling instant when wall penetration or exit is detected.




Figure 71: Time Response of Oscillations with 1 kHz Sample Rate and 2650 counts/rev




Figure 72: Phase Plane Plot of Oscillations with 1 kHz Sample Rate and 2650 counts/rev

11.2 Impulse Damping

As Gillespie showed (1996), when a manipulandum crosses into or out of a virtual wall, the asynchrony between the actual crossing time and the time at which the controller updates the commanded torque will lead to a finite torque error. In systems with relatively low resolution displacement sensors (e.g., optical encoders) this can lead to oscillations that may only occur between two adjacent encoder values. These oscillations are simpler than those for a continuous resolution displacement sensor in two ways. First, since the position (and thus the commanded spring torque) does not get updated when the manipulandum is within the wall, the ZOH energy leak is not relevant, and one must only address the threshold crossing asynchrony energy leak (as shown in Chapter 9). Second, rather than the torque-time integral error being a continuous function that requires knowledge of the system model or accurate trajectory information to determine, it is a constant function of the error between actual threshold crossing time and the next sample instant:

$$\int T_{error} dt = K\Delta\theta \frac{t_s}{2},$$

where T_{error} is the torque error equal to the spring constant, K , times the displacement resolution quantum, $d\theta$, and the average error time is assumed to be half the sampling interval, t_s .

Impulse damping involves the application of a torque impulse at the beginning of the next sample interval after a threshold crossing is detected in order to cancel the estimated average torque error due to the asynchrony. The magnitude is simply the torque-time integral for the error, scaled by the ratio of the sampling period to the impulse period:

$$T_I = K\Delta\theta \frac{t_s/2}{t_I},$$

where T_I is the corrective impulse torque.

An early implementation of this scheme applied the correction torque during the entire next sampling period, but this lengthy application of the correction torque sometimes continued past the point that the manipulandum reversed direction, and the “correction” torque actually accelerated the manipulandum for the latter portion of the sample period. Practical implementation of this technique will likely require a sub-controller running 10-40 times faster than the main sample rate to more accurately detect crossing times and to time the pulse. The initial simulation results presented here assumed a fast sub-controller to command the corrective pulse, but did not take advantage of this faster rate for more accurate threshold crossing detection. Fast, large magnitude pulses may also challenge the response times of the haptic amplifier and actuator.

Figure 71 illustrated simulated limit cycle oscillations for a manipulandum in contact with a virtual wall with a sample rate of 1 kHz and an encoder resolution of 2650 counts/revolution. Figure 73 shows a similar plot for a simulation with impulse damping occurring at the beginning of the next interval after detection of a threshold crossing. The damping impulse period was 1/40 of the main sample period. Notice the large oscillation to the right of the plot, amidst smaller oscillations. This is caused by a “worst case” situation where the threshold crossing occurred just after a sample instant, and the torque-time integral error was a full $K\Delta\theta t_s$, rather than $K\Delta\theta t_s/2$, as assumed for the

impulse damping compensation. These worst-case pulses occurred occasionally, and would most likely be perceptible by the human grasping the manipulandum.




Figure 73: Simulated Limit Cycle Oscillations at 1 kHz Sample Rate, with 2650 counts/revolution Encoder, with Impulse Damping Applied

Figure 74 contains a phase plane plot for the Impulse Damping example. The effects of the impulses appear as vertical lines at the juncture between the inside-the-wall arcs and the outside-the-wall arcs. The impulses “knock down” the manipulandum velocity, which limits the displacement excursion, but notice the phase plane plot becomes more slender (less wide with respect to height), which is a good indication that impulse damping increases the frequency of the oscillations (as one can also see in the time response plot). An ideal mitigation scheme would decrease the size of the oscillations, but maintain the same aspect ratio for the trajectory (ratio of velocity magnitude to displacement magnitude, which is indication of frequency).




Figure 74: Phase Plane Plot for Impulse Damping Example

Stochastic variations in the bounces occurring with an impulse damping scheme can limit its utility. For a system with position updates occurring at a set rate (e.g., 1 kHz) and impulse damping occurring at the beginning of every 1 ms interval following a threshold crossing, oscillation magnitude will be significantly reduced, but the occasional “worst case” occurrence of the manipulandum “sneaking across” a threshold just after a sample instant, resulting in a torque error for the maximum possible time, will result in an unusually large bounce. If the main bounce frequency is 270 Hz, and these errant bounces occur every 10 bounces, the errant bounce rate will be 27 Hz, and these larger bounces will be easily perceptible. In addition to the occasional errant bounce, note that the smaller bounces increase in frequency to approximately 270 Hz, compared to approximately 80 Hz for the nominal bounces.

One potential refinement to the technique is to keep the main loop rate constant (e.g., ZOH commands at 1 kHz), but to increase the frequency of threshold detection to some

multiple (e.g., 20 kHz). The lower main loop rate keeps the oscillation frequency down, while using more precise knowledge of the threshold crossing times to adjust the gains of the damping impulses to more closely approximate the actual torque-time integral error.

The Impulse Damping technique has a theoretical problem that limits its effectiveness. It seeks to correct the torque-time integral, applying a compensating impulse with the same torque-time product as the preceding error. But we have seen that energy leakage corresponds to the integral of the product of torque and velocity over time:

$$E_{leak} = \int_{t_1}^{t_2} T_{error}(t) \dot{\theta}(t) dt$$

Unfortunately, since we lack velocity information, this integral cannot be calculated. The lack of a velocity term in the calculation of our compensation pulse means that the “damping” portion of the “Impulse Damping” name is more a wish than reality. If we knew the oscillation magnitude and zero crossing times and assumed a sinusoidal signal, we could estimate velocity, but we do not have adequate displacement information, as this is the problem we are attempting to solve! The second mitigation technique, Phase Estimation Damping, attempts to use threshold crossing information despite the absence of displacement and velocity information.

11.3 Phase Estimation Damping

The impulse damping approach seeks to apply corrective torque pulses to compensate for asynchronous threshold crossing errors. Impulse damping implemented with a single compensation pulse for every sample interval during which a zero crossing

occurs has limitations due to the fact that it assumes “average case” threshold crossing asynchrony, with threshold crossings occurring on average in the middle of the sample interval. Worst case threshold crossings (right after a sample instant) can cause relatively large uncompensated torque errors leading to large oscillatory excursions. Increasing the threshold detection rate while leaving the main loop rate constant in order to tune the magnitude of the damping impulses may lead to some improvement, but oscillations will remain. An alternative approach is to attempt to provide damping torques, which are essentially torques out of phase with displacement, regardless of the source of oscillation energy. The objective is to reduce oscillation magnitude without necessarily increasing oscillation frequency.

For fairly regular limit cycle oscillation waveforms, if one lacks position and velocity information but possesses accurate information as to zero-crossing times, one can estimate the phase of the oscillations and produce forces (torques) out of phase with velocity, effectively applying damping. This damping could be produced by sinusoidally varying torques leading the estimated displacement trajectory by 90 degrees of phase angle. This maximizes the damping torque at maximum velocity, which also happens to be at the threshold crossing where knowledge of manipulandum position is most accurate (in fact, the only time it is truly known). Damping torques taper off sinusoidally as phase angle moves away from the zero crossing and towards the direction reversal point. Phase estimation errors can result in torques being applied to accelerate the manipulandum rather than decelerate it – luckily these erroneous torques would likely occur near the direction reversal point, where damping torques would be lowest. The damping

controller would continually update the zero-crossing interval value, making adjustments in the damping waveform to accommodate any changes in oscillation frequency.

11.3.1 INITIAL PHASE ESTIMATION DAMPING RESULTS

Figure 75 contains the time response plot, and Figure 76 contains the phase plane plot for an initial attempt at Phase Estimation Damping. The technique, as currently implemented, does not have much value. Oscillation magnitudes vary widely, and in this example Phase Estimation Damping does not appear to provide any reliable benefit.




Figure 75: Time Response Plot for Initial Attempt at Phase Estimation Damping

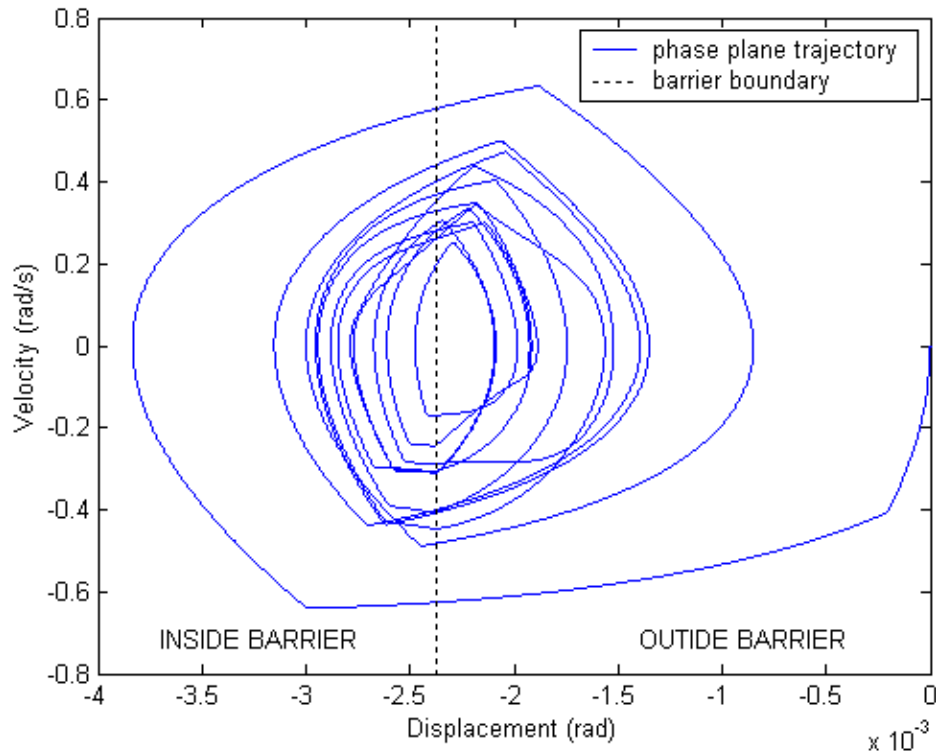


Figure 76: Phase Plane Plot for Initial Attempt at Phase Estimation Damping

One limitation of the current approach is that it only measures the barrier penetration times, and assumes that the following wave period will have a symmetrical sinusoidal shape. Since different control laws operate on either side of the barrier, this assumption is not a good one. The phase plane plot for the nominal case, in Figure 72, with its different shaped arcs for the inside-wall controller and outside-wall controller, exemplifies this point. Changing bias torques also destroy the symmetry of the waveform about the wall boundary. Measuring the exits as well as the penetrations, resetting the phase estimate upon every penetration and exit, and estimating each out-of-wall time interval based on the last out-of-wall interval, and each inside-wall interval based on the last inside-wall interval, may provide better results by improving the accuracy of the phase estimate. More intelligent phase estimation schemes, perhaps using Kalman filtering, could provide valuable benefits.

The current Phase Estimation Damping scheme attempts to apply a damping vector for a full 360 phase degrees after each penetration. Rather than simply updating at each barrier crossing and estimating for 180 phase degrees in an effort to provide a damping vector up until the next barrier crossing, a more conservative approach would only apply a damping vector for 90 phase degrees after a barrier crossing to minimize the chance that a phase error would result in a torque being applied for a time that instills energy rather than dissipates it (e.g., incorrectly anticipating the point at which knob direction reverses). This approach has similarities to work by Cham et al. (2001), who were working to *maximize* the energy transferred to an oscillating leg of a hopping robot.

11.4 Implications for Multi-Degree-of-Freedom Systems

The suggested approach of including a fast sub-controller “ganglion” in a system is relatively straightforward for a haptic knob, where kinematic and inverse Jacobian calculations are not needed. In a typical multi-degree-of-freedom haptic system, the kinematic calculations to convert joint-space displacements to Cartesian displacements and the inverse Jacobian calculations to convert Cartesian forces and torques to joint-space torques take place at the haptic servo rate. Kinematic and inverse Jacobian calculations could be difficult and inefficient to “hard-wire” into custom digital logic, which might lead one to question the value of the suggested approach for anything but single degree-of-freedom systems or systems with completely decoupled degrees of freedom (such as a haptic joystick with a spherical five-bar linkage).

Kinematics and inverse Jacobian coefficients change with changes in the kinematic configuration of the haptic structure (as different parts of the device move relative to each other). At velocities common in haptic feedback applications, these coefficients are not

likely to change much during very small intervals such as the 21 microseconds in between 48 kHz PWM updates or even during the 1 millisecond between 1 kHz haptic loop updates. This encourages consideration of a “joint-space proxy” approach, with a spring (or even a spring-damper, assuming high-enough displacement resolution for an adequate velocity signal) calculated at a very high rate (such as the PWM rate) and the kinematics and inverse Jacobian calculated at the slower main haptic loop rate.

Another factor to consider is that with costs of haptic devices scaling exponentially with the number of degrees of freedom, higher degree-of-freedom device designers are less likely to feel pressure to make do with lower resolution displacement sensors, making the application of the suggested approach to multi-degree-of-freedom systems moot.

Chapter 12 : Conclusions

The initial goal of this dissertation was to understand the relative roles of displacement quantization and the zero-order hold in creating limit cycle oscillations upon contact with a virtual barrier. Contrary to initial expectations, the two factors are not independent. Displacement quantization acts to magnify ZOH-induced energy leaks, and does not by itself generate energy leaks. As a prerequisite to simulation of a haptic knob with displacement quantization and ZOH, extensive system identification studies were conducted of the human hand holding a haptic knob in various grasp postures. Results showed that for light to moderate grip forces, a second-order model adequately approximates the dynamic behavior of the hand/knob combination, and that for strong grip forces, a fourth order model offers a plausible explanation of behavior. For the second-order model estimates, both stiffness and damping increase monotonically with grip force.

Displacement quantization exacerbates limit cycle oscillations primarily by instilling more energy upon exit from a virtual barrier than a discrete controller without displacement quantization would instill due to the previously recognized ZOH leak effects. The errant torque pulse for a quantized displacement system typically has a large

magnitude and occurs at precisely the worst moment, when velocity out of the barrier is the greatest and the most work can be done on the knob. Though the time integral of torque error does not rigorously describe the energy leak, multiplying the average ZOH lag (half the sample period) by the quantization-induced exit torque error offers an excellent prediction of the sensitivity of limit cycle magnitude to variations in sample rate and displacement quantization. The work on the effects of displacement quantization applies not only to quantization due to optical encoders or other discrete sensors, but also to quantization from A/D, D/A, PWM, or other sources.

The frequency profiles of human sensitivity to tactile and auditory oscillations should motivate haptic system designers to attempt to reduce limit cycle oscillations without increasing limit cycle frequency. Increasing sample rate can easily increase limit cycle oscillations frequencies to the point that limit cycles can be more easily heard than felt. In fact, in some instances with high sample rates and audible limit cycles, designers should consider lowering sample rates. Increasing sample rate decreases oscillation magnitude but increases oscillation frequency; however, increasing displacement resolution reduces limit cycle magnitude without much effect on frequency. This makes increasing displacement resolution rather than increasing sample rate the preferred method of avoiding noxious limit cycles in haptic systems. This result differs from conclusions one may have drawn from previous literature on haptic limit cycles that since ZOH lags are the root of the problem, increasing sample rates is an obvious way to treat this problem.

Two tools for estimating limit cycle behavior were introduced. The first, a simple approximation of limit cycle magnitude obtained by multiplying the displacement

quantization, sample delay, and a constant, generates useful results if one has a known case from which to obtain the constant. Given a case with known sample rate, displacement resolution, and limit cycle magnitude, one can easily estimate the effect on oscillation magnitude of changing sample rate or displacement resolution. The second technique, describing function analysis, can give estimates of both oscillation magnitude and oscillation frequency, and does not require a known case (though it does require an accurate model of the hand and haptic knob).

Two methods for mitigating limit cycle oscillations despite the use of coarse encoders were explored, with limited success. Potential improvements to both methods were identified, and each may be worthy of further work. Even if methods for quenching limit cycles with coarse encoders were developed, other adverse perceptual effects such as detectable torque steps in a virtual barrier will limit how low a resolution is acceptable. The difficulty of mitigating limit cycles with coarse encoders and the superiority of quenching limit cycles with increased displacement resolution rather than increased sample rate provides motivation for R&D to develop inexpensive high-resolution encoders, if the development costs can be justified by later savings in manufacturing costs or the enabling of new applications with requirements for high haptic fidelity and low cost.

Appendix A : Motor Calibration

Kearney and Hunter caution against the use of motor torque commands as a substitute for force or torque data, stating that the current-torque relationships of motors have dynamics, nonlinearities, and time-varying characteristics that can lead to serious system identification errors. They assert that independent measurement of torque is essential (Kearney and Hunter, 1990). This guidance encouraged the inclusion of torque sensing capability in the HREF system. Unfortunately, the torque load cell exhibited significant vulnerability to off-axis forces and torques. Torque measurements from this sensor contained nonlinearities and other artifacts that prevented the use of the sensor data. This left the record of commanded motor torque as the most attractive method of estimating torque. With Kearney and Hunter's caution in mind, the HREF system and its Maxon RE025 motor were carefully characterized for static current-torque behavior, dynamic friction, and transient response.

A.1 Torque Constant and Dynamic Friction

Motor calibration was accomplished using a Vibrac Model 1038 motorized dynamometer from the Vibrac Corporation, Amherst, NH. The Model 1038 was connected to a Model 6610 test head that carried the dynamometer actuator and Model TQ-10 torque sensor. The entire HREF device (without a knob) was mounted vertically atop the dynamometer test head. The motor shaft of the Maxon RE025 motor in HREF, and the shaft of the dynamometer, were both oriented vertically, and were concentrically aligned as close as possible. The test head included two flexible couplings – one between the RE025 and the torque sensor, and the other between the torque sensor

and the dynamometer drive actuator – to minimize any effects of misalignment and off-axis loads. The dynamometer saw the Maxon motor, with its bearings and brushes, as well as the two HREF bearings, as its load (recall that the Maxon RE025 motor in HREF is suspended by its shaft on two additional bearings to allow the motor housing to rotate against a reactive torque arm).

For a given drive current into the Maxon RE025 motor, the dynamometer slowly drove the RE025 shaft clockwise for approximately three rotations and then counterclockwise for approximately three rotations. Torque and shaft angle were sampled every 100 ms and recorded. Clockwise and counterclockwise torque measurements were averaged separately, and the difference taken to indicate the dynamic friction torque. The two figures below show plotted data and fitted lines for motor torque and friction as a function of drive current. Figure 77 illustrates behavior for very small motor torques (less than 2 mNm). Figure 78 illustrates behavior for the range of torques relevant to the current research (0-30 mNm). These results show extremely linear torque response and constant dynamic friction across both ranges. The torque constants obtained for the two ranges (26.4 mNm for the lower range and 28.4 mNm for the higher range) differ by 7%, and are both substantially higher than the 23.5 mNm/A quoted in Maxon technical literature for the RE025-118752 motor (Maxon, 1999).




Figure 77: Motor Torque and Dynamic Friction as a Function of Drive Current for Small Torques (less than 2 mNm)




Figure 78: Motor Torque and Dynamic Friction as a Function of Drive Current for the Relevant Operating Range (0-30 mNm)

A.2 Transient Response

Use of the commanded motor torque as an indication of actual torque presumes that the motor output comes close enough to the commanded torque to avoid introducing unwanted artifact into the results. In addition to characterizing the current-torque gain and friction losses, one must ask whether the torque output of the motor actually tracks the commanded output well with time. In particular, given a step or pulse command, does the motor current ramp so quickly and settle so quickly that these nonidealities may be disregarded for the purposes of the current investigation?

Figure 79 shows the transient response of the Trust Automation TA115 amplifier driving the Maxon RE025 motor in HREF. The waveform was obtained from the current sensing circuit inside the TA115 amplifier. This sensing circuit is composed of a two-stage amplifier with a differential input around either end of a current sense resistor, and a second stage for scaling and filtering. Since the gain for this signal is set for use as a feedback signal, positive output current is represented by a negative voltage. In the plot below, the asymptotic signal value of approximately -2 V corresponds to an input voltage command of 2 V, corresponding to a current command of 0.4 A. Each time division in the plot below denotes 100 microseconds, which is also the sample interval used for the system identification experiments. One can see that the amplifier-motor combination settles in less than 200 microseconds – much too fast to significantly disturb the system identification results.




Figure 79: Transient Response of the Amplifier and Motor in HREF

Appendix B : Grip Force Load Cell Calibration

The grip load cell was calibrated using lead weights the masses of which were verified using an Ohaus triple-beam balance. During load cell calibration, the weights were hung at one end of a horizontal beam that was hinged at the other end. The grip-force-sensing knob assembly rested under the center of the beam, receiving twice the weight of the test masses due to the lever action of the beam. Where the beam rested on the knob, it was ground to a 'knife edge' to allow precise location of the point of force application on the knob surface. Figure 80(a) shows the calibration beam resting across the knob, applying a force directly over the load cell embedded in the knob. The beam rests on the 'shell' part of the knob, which is attached to the main knob body with a flexure. The flexure acts as a fulcrum, and any forces applied closer to the fulcrum than the load cell will be underestimated, while forces applied farther from the fulcrum will be overestimated. This applies both for calibration and for tests with human subjects where finger position on the knob may vary. Figure 80(b) shows a side view of the knob, with callouts indicating the flexure and the load cell.




Figure 80: (a) Knob with calibration beam; (b) Closeup of knob

Figure 81 shows the results of the grip load cell calibration. Data points for test forces applied directly over the center of the load cell (the intended point of finger force application) appear as dots. This nominal point of force application lies 6.35 mm (0.25") from the tip of the knob. The solid line represents a least-squares fit to this data, and produced the slope and grip cell coefficient for the "nominal case" in the plot. The grip cell coefficient is simply the slope inverted and scaled to the proper units for use in converting sensor voltage to grip force during trials and later analysis. In order to estimate potential errors from fingers squeezing at a point on the knob closer to or farther away from its end than 6.35 mm (1/4"), data was collected for two "deviant" cases – one with forces applied 3.18 mm (1/8") from the end of the knob, and the other 9.53 mm (3/8") from the end of the knob.

The two dashed lines in Figure 81 show predicted calibration curves for a flexure acting as a simple rigid hinge. The triangles show calibration data for forces applied 3.18 mm from the end of the knob. The squares show calibration data for forces applied 9.53 mm from the end of the knob. These data points lie significantly outside the dashed lines showing predictions for the simple rigid hinge assumption. Clearly, the 1.9 mm (0.075") of free flexure acts with the other parts of the knob assembly in a way that violates the simple rigid hinge assumption, worsening the deviance of the estimated load for forces applied away from the nominal location. The two dotted lines show least-squares fits for the two deviant cases. The 9.53 mm case has a grip cell coefficient nearly twice that of the 3.18 mm case. Estimates of the same force applied to these two extreme locations would differ by nearly a factor of two. Note that the extremes used for the test do not necessarily correspond to actual variation in the finger placement of subjects.

Zero-mean stochastic variations of finger placement around the nominal position would not significantly affect the results of this work. Systematic errors (subjects consistently placing fingers to one side of the nominal position) would skew the correspondence between grip force and system identification parameters. This skew would not reduce the validity of the qualitative results or effects observed, but could harm the validity of system parameter estimates based on specific grip force assumptions. If a designer using these results assumed a specific grip force for a particular scenario and desired human finger dynamic parameter estimates for that scenario, systematic errors in finger placement during the present study would skew the estimates he or she obtained.




Figure 81: Grip Force Load Cell Calibration Plot (with calibration masses doubled to account for lever arm in experimental setup).

Appendix C : Apparatus Dynamics

In order to obtain an estimate of human hand dynamics in a knob grasp, the dynamics of the apparatus must be subtracted from the overall dynamics estimate obtained during an experiment. Apparatus dynamics were estimated from 3D CAD data (for moment of inertia), empirical testing with the apparatus, and selected hand calculations. Empirical testing yielded reliable data for the bare motor shaft setup, small (17.8 mm diameter) knob setups, and medium (29.9 mm diameter) setups. Large (45 mm diameter) knobs were constructed, but their high inertia prevented reliable measurements, so only small and medium knob setups were used for human grasp identification experiments.

Empirical measurements of moment of inertia, damping, and stiffness were taken with sets of trials of pulses of the same six magnitudes as those used for human grasp identification experiments. Damping estimates for all the apparatus setups used in the human grasp identification experiments were more than an order of magnitude lower than the lowest damping estimates obtained in the grasp experiments (which occurred for the lightest grip force on the small knob). Stiffness estimates for apparatus setups were more than two orders of magnitude below the lowest stiffness measured for the grasp experiments (also for the lightest grip force on the small knob). Based on these results, inherent damping and stiffness in the apparatus was regarded as negligible for the analysis of the human grasp identification experiments. Moments of inertia represent the opposite situation; apparatus inertias are quite high compared to the inertias of human

fingers in various knob grasps. This requires consideration of the apparatus inertias during analysis of the human grasp experiments.

Table 9 contains estimates of moments of inertia obtained from within the SolidWorks 3D CAD software application (Schena, 2001; SolidWorks Corporation, 2001). The moments of inertia in Table 9 do not include the inertia of the motor rotor, shaft, sensor hub, and encoder disk, which were determined empirically. The "proof setups" contained solid knob parts with no grip force sensor. The "experimental setups" contained the specially-machined knob core with a grip force sensor. The three calculated moment of inertia estimates useful for comparison with empirical estimates appear in bold at the bottom of Table 9. The moment of inertia of the rotor, sensor hub, and encoder disk must be added to these estimates to represent the total moment of inertia of each apparatus configuration. A hand calculation for the small solid proof knob using the formula for the moment of inertia of a cylinder, $J = \frac{1}{2}MR^2$, yielded an estimate of 7.75 gm-cm², almost identical to the 3D CAD estimate of 7.74 gm-cm².

Table 9: Moments of Inertia Calculated in SolidWorks 3D CAD Program, in Units of gm-cm² (Schena, 2001).

	Small Proof Setup	Small Experimental Setup	Medium Proof Setup	Medium Experimental Setup	Large Proof Setup	Large Experimental Setup
17.8 mm knob (solid)	7.74		7.74		7.74	
17.8 mm knob (w/grip sensor)		7.55		7.55		7.55
29.9 mm knob cap (one piece)			11.92			
45 mm knob cap (one piece)					66.68	
29.9 mm knob cap (split)				N/A		
45 mm knob cap (split)						N/A
knob cap mounting screws (4)			0.65	0.65	0.65	0.65
TOTALS	7.74	7.55	20.31	N/A	75.07	N/A

* Totals do not include moment of inertia of motor rotor, shaft, encoder disk, and sensor hub.

** Calculated moment of inertia for knob core with grip force sensor does not include two small overload screws, a set screw, and four small, flexible load cell wires.

Table 10 contains empirical estimates of the moment of inertia for the six apparatus configurations. A similar moment of inertia estimate was also obtained for the motor with a bare shaft, representing the inertia of the rotor, shaft, sensor hub, and glass encoder disk. This estimate is subtracted from the overall estimates in Table 10 to provide knob inertia estimates (the three figures in bold) for direct comparison to the three values in bold in Table 9. For the two small knob configurations, the empirical estimates come within 5% of the values calculated in the CAD program. The medium proof setup has an error of 10.8%, and the large proof setup has an error of over 30%, showing a trend towards unreliable estimates for higher moments of inertia. The underlined empirical estimates for the small experimental setup and the medium experimental setup were chosen as representations of the apparatus moments of inertia for use during analysis of the human grasp identification experiments.

Table 10 Empirical Estimates of Moment of Inertia

	Small Proof Setup	Small Experimental Setup	Medium Proof Setup	Medium Experimental Setup	Large Proof Setup	Large Experimental Setup
Empirical estimate	20.2	<u>19.8</u>	34.5	<u>32.7</u>	64.2	68.7
minus empirical estimate of motor rotor, shaft, sensor hub, and disk	-12.0	-12.0	-12.0	-12.0	-12.0	-12.0
TOTALS	8.14	7.8	22.5	20.6	52.2	56.7
Error	5.2%	3.2%	10.8%		-30.5%	

Appendix D : Raw Data for Empirical Trials




Figure 82: Raw Data for Trial at 455 Hz and 256 counts/revolution




Figure 83: Raw Data for Trial at 1000 Hz and 256 counts/revolution




Figure 84: Raw Data for Trial at 2 kHz and 256 counts/revolution




Figure 85: Raw Data for Trial at 4 kHz and 256 counts/revolution




Figure 86: Raw Data for Trial at 455 Hz and 512 counts/revolution




Figure 87: Raw Data for Trial at 1 kHz and 512 counts/revolution




Figure 88: Raw Data for Trial at 2 kHz and 512 counts/revolution




Figure 89: Raw Data for Trial at 5 kHz and 512 counts/revolution




Figure 90: Raw Data for Trial at 455 Hz and 1024 counts/revolution




Figure 91: Raw Data for Trial at 1 kHz and 1024 counts/revolution




Figure 92: Raw Data for Trial at 2 kHz and 1024 counts/revolution




Figure 93: Raw Data for Trial at 5 kHz and 1024 counts/revolution




Figure 94: Raw Data for Trial at 455 Hz and 2048 counts/revolution




Figure 95: Raw Data for Trial at 1 kHz and 2048 counts/revolution




Figure 96: Raw Data for Trial at 2 kHz and 2048 counts/revolution




Figure 97: Raw Data for Trial at 5 kHz and 2048 counts/revolution

References

- R. J. Adams and B. Hannaford, '[A Two-Port Framework for the Design of Unconditionally Stable Haptic Interfaces](#),' Proceedings of IROS 98, pp. 1254-59, Victoria, B.C., Canada, Nov. 1998.
- R. J. Adams and B. Hannaford, "Stable Haptic Interaction with Virtual Environments," *IEEE Transactions on Robotics and Automation*, Vol. 15, No. 3, January 1999, pp. 465-474.
- R. J. Adams and B. Hannaford, '[Control Law Design for Haptic Interfaces to Virtual Reality](#),' In Press: IEEE Trans. Control Systems Technology, June 2001.
- R. Adams, M. Moreyra, B. Hannaford, "Stability And Performance of Haptic Displays: Theory and Experiments," *Proceedings of the ASME IMECE*, Haptics Workshop, Nov 1998, pp. 227-234.
- R. J. Adams, D. Klowden, and B. Hannaford, '[Stable Haptic Interaction Using the Excalibur Force Display](#),' *IEEE Intl. Conf. on Robotics and Automation*, San Francisco, CA, April 2000.
- G. C. Agarwal and G. L. Gottlieb, "Compliance of the human ankle joint," *ASME Journal of Biomechanical Engineering*, Vol. 99, pp. 166-170, 1977.
- R. J. Anderson and M. W. Spong, "Asymptotic Stability for Force Reflecting Teleoperators with Time Delay," *Int. Journal of Robotics Research*, Vol. 11, No. 2, 1992, pp. 135-149.
- S. Ashley, "Simplifying Controls," *Automotive Engineering International*, March 2001.
- J. D. Becker and C. D. Mote Jr., "Identification of a frequency response model of joint rotation," *Journal of Biomechanical Engineering*, Vol. 112, pp. 1-8, 1990.
- A. S. Birch, "Experimental Determination of the Viscoelastic Properties of the Human Fingerpad," S.M. thesis, MIT, Cambridge, MA, 1998.
- P. Bogacki and L. F. Shampine, "A 3(2) pair of Runge-Kutta formulas," *Appl. Math. Letters*, Vol. 2, 1989, pp. 1-9.

- J. G. Cham, J. Karpick, J. E. Clark, and M. R. Cutkosky, "Stride Period Adaptation for a Biomimetic Running Hexapod" *10th International Symposium of Robotics Research*, Lorne, Victoria, Australia, November 9th - 12th, 2001.
- J. E. Colgate, "Robust Impedance Shaping Telemanipulation," *IEEE Trans. Robotics and Automation*, Vol. 9, No. 4, 1993, pp. 374-384.
- J. E. Colgate, personal communication, September 2001.
- J. E. Colgate, et al., "Implementation of Stiff Virtual Walls in Force-Reflecting Interfaces," *Proc. IEEE Virtual Reality Annual Int. Symp.*, Seattle, WA, 1993, pp. 202-208.
- J. E. Colgate, M. C. Stanley, and J. M. Brown, "Issues in the Haptic Display of Tool Use," *Proc. IEEE/RSJ Int. Conf. On Intelligent Robots and Systems*, Pittsburgh, PA, 1995, pp. 140-145.
- M. Cutkosky, personal communication, information available at <http://www-cdr.stanford.edu/Touch/>, October, 2001.
- R. Crowninshield, M. H. Pope, R. Johnson, and R. Miller, "The impedance of the human knee," *Journal of Biomechanics*, Vol. 9, No. 8, pp. 529-535, 1976.
- A. A. H. Damen, "Modern Control Theory: Lecture Notes for Course 5N050," <http://www.er.ele.tue.nl/ADamen/modcon/total.pdf>, Measurement and Control Group, Department of Electrical Engineering, Eindhoven University of Technology, the Netherlands, June, 2001.
- J. M. Dolan, M. B. Friedman, and M. L. Nagurka, "Dynamic and Loaded Impedance Components in the Maintenance of Human Arm Posture," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 23, No. 3, pp. 698-709, May/June 1993.
- R. E. Ellis, N. Sarkar, and M. A. Jenkins, "Numerical methods for the haptic presentation of contact: Theory, simulations, and experiments," *Proceedings of the ASME Dynamic Systems and Control Division*, DSC-Vol. 58, pages 413-420, 1996.
- R. E. Ellis, N. Sarkar, and M. A. Jenkins, "Numerical Methods for the Force Reflection of Contact," *ASME Trans. On Dynamic Systems, Modeling, and Control* 119(4):768-774, 1997.
- A. Gelb and W. E. Vander Velde, "Multiple-Input Describing Functions and Nonlinear System Design," McGraw-Hill, 1968.

- B. Gillespie, P. Tang, and C. Hasser, "Cancellation of Feedthrough Dynamics Using a Force-Reflecting Joystick," *ASME International Mechanical Engineering Conference and Exposition*, Nashville, TN, November 1999.
- B. Gillespie and M. Cutkosky. "Stable User-Specific Rendering of the Virtual Wall," *Proceedings of the ASME International Mechanical Engineering Conference and Exposition*, DSC-Vol. 58, Atlanta, GA, Nov 17-22, 1996. pp. 397-406.
- M. D. Good, J. B. Munson, "Torque motor as a tactile feedback device in a computer system," U.S. Patent No. 5,185,561, February 9, 1993.
- M. Gotfrit, Ed., "Zen & the Art of Audio: A web based audio primer," <http://www.sfu.ca/sca/Manuals/ZAAPf/new%20page.html>, Simon Fraser University, Burnaby, B.C., Canada, 2001.
- R. J. Gulati, "Determination of mechanical properties of the human fingerpad, in vivo, using a tactile stimulator," MS thesis, Boston University, Boston, MA, 1995.
- R. J. Gulati and M. A. Srinivasan, Human fingerpad under indentation I: static and dynamic force response, *Proceedings of the 1995 Bioengineering Conference*, Eds: R. M.
- A. Z. Hajian, "A Characterization of the Mechanical Impedance of Human Hands," Ph.D. dissertation, Harvard University, September, 1997.
- A. Z. Hajian and R. D. Howe, "Identification of the mechanical impedance at the human finger tip," *ASME Journal of Biomechanical Engineering* 119(1):109-114, Feb. 1997. Also presented at the International Mechanical Engineering Congress, American Society of Mechanical Engineers, Chicago, IL, November 1994, *Proceedings* ed. C. J. Radcliffe, DSC-vol. 55-1, p. 319-327.
- A. Z. Hajian, D. S. Sanchez, and R. D. Howe, "Drum Roll: Increasing Bandwidth Through Passive Impedance Modulation," *IEEE Int. Conf. On Robotics and Automation*, Albuquerque, NM, April 1997.
- B. Hannaford, "A Design Framework for Teleoperators with Kinesthetic Feedback," *IEEE Trans. Robotics and Automation*, Vol. 5, No. 4, 1989, pp. 426-434.
- B. Hannaford, personal communication, September 2001.
- B. Hannaford, J.H. Ryu, Y.S. Kim, '[Stable Control of Haptics.](#)' In "Touch in Virtual Environments: Proceedings USC Workshop on Haptic Interfaces, Feb 23 2001", Margret McLaughlin, Ed., Prentice Hall, 2001.
- B. Hannaford, J.H Ryu, '[Time Domain Passivity Control of Haptic Interfaces.](#)' Proc. IEEE Intl. Conference on Robotics and Automation, Seoul, May 2001.

- V. Hayward, J. Choksi, G. Lanvin, and C. Ramstein, "Design and multi-objective optimization of a linkage for a haptic device," In *Proceedings of the 4th Workshop on Advances in Robot Kinematics*, Ljubljana, Slovenia, 1994.
- N. Hogan, "Mechanical Impedance of Single- and Multi-Articular Systems," *Multiple Muscle Systems: Biomechanics and Movement Organization*, Chapter 9, J. M. Winters and S. L-Y. Woo (ed.), Springer-Verlag, New York, 1990.
- I. W. Hunter and R. E. Kearney, "Dynamics of Human Ankle Stiffness: Variation with Mean Ankle Torque," *Journal of Biomechanics*, Vol. 15, No. 10, pp. 747-752, 1982.
- I. W. Hunter and R. E. Kearney, "Invariance of ankle dynamic stiffness during fatiguing muscle contractions," *Journal of Biomechanics*, Vol. 16, No. 12, pp. 985-991, 1983.
- International Standards Organization. 1961. *Normal Equal-Loudness Contours for Pure Tones and Normal Threshold of Hearing under Free Field Listening Conditions R226*. ISO: New York, NY.
- R. Johansson and G. Westling, "Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects," *Experimental Brain Research*, Vol. 56, pp. 550-564, 1984.
- E. I. Jury, "Theory and Application of the Z-Transform Method," J. Wiley, New York, 1964.
- S. Karason and M. A. Srinivasan, "Passive human grasp control of an active instrumented object," *Proceedings of the International Mechanical Engineering Congress and Exposition*, ASME, Chicago, IL, DSC Vol. 57-2, pp. 641-647, November, 1995.
- Katz, J. (1985). *Handbook of Clinical Audiology*, 3rd ed. (Williams and Wilkins, Baltimore).
- R. E. Kearney and I. W. Hunter, "Dynamics of human ankle stiffness: Variation with displacement amplitude," *Journal of Biomechanics*, Vol. 15, No. 10, pp. 753-756, 1982.
- R. E. Kearney and I. W. Hunter, "System Identification of Human Joint Dynamics," *Critical Reviews in Biomedical Engineering*, Vol. 18, Issue 1, pp. 55-87, 1990.
- R. E. Kearney, R. B. Stein, and L. Parameswaran, "Identification of intrinsic and reflex contributions to human ankle stiffness dynamics," *IEEE Transactions on Biomedical Engineering*, Vol. 44, No. 6, pp. 493-504, 1997.

- L. Love and W. Book, "Contact stability analysis of virtual walls," *Proceedings of the International Mechanical Engineering Congress and Exposition*, ASME, Chicago, IL, DSC Vol. 57-2, pp. 689-694, November, 1995.
- K. E. MacLean, C. J. Hasser, and L. Chu, "Driving with Programmable Haptic Feedback: Design Scenarios and Contextual Evaluation," *in review*.
- Maxon Precision Motors, Inc., *1999 Main Catalogue*, 838 Mitten Road, Burlingame, CA.
- MicroE Systems Corporation, corporate literature, <http://www.microesys.com>, Natick, MA, 2001.
- B. E. Miller, J. E. Colgate, and R. A. Freeman, "Passive Implementation for a Class of Static Nonlinear Environments in Haptic Display," *Proc. 1999 IEEE Int. Conf. On Robotics and Automation*, Detroit, Michigan, May 1999, pp. 2937-2942.
- B. E. Miller, J. E. Colgate, and R. A. Freeman, "Computational Delay and Free Mode Environment Design for Haptic Display," *Proc. ASME Dynamic Systems and Control Division, DSC-5B-4*, ASME, 1999.
- T. E. Milner and D. W. Franklin, "Two-dimensional endpoint stiffness of human fingers for flexor and extensor loads," *Proceedings of the International Mechanical Engineering Congress and Exposition*, ASME, Chicago, IL, DSC Vol. 57-2, pp. 649-656, November, 1995.
- T. E. Milner and D. W. Franklin, "Characterization of Multijoint Finger Stiffness: Dependence on Finger Posture and Force Direction," *IEEE Transactions on Biomedical Engineering*, Vol. 45, No. 11, pp. 1363-1375, 1998.
- F. Mougnet and V. Hayward, "Limit Cycle Characterization, Existence, and Quenching In the Control of a High Performance Hydraulic Actuator," *Proc. IEEE International Conference on Robotics and Automation*, Vol. 3, pp. 2218-2223, 1995.
- G. Niemeyer, personal communication, October, 2001.
- G. Niemeyer and J. J. Slotine, "Stable Adaptive Teleoperation," *IEEE Journal of Oceanic Engineering*, Vol. 16, pp. 152-162, 1991.
- D.T.V. Pawluk, and R.D. Howe, [Dynamic Lumped Element Response of the Human Fingerpad](#), *ASME Journal of Biomechanical Engineering* 121(2):178- 184, April 1999.
- D.T.V. Pawluk, and R.D. Howe, [Dynamic contact of the human fingerpad against a flat surface](#), *ASME Journal of Biomechanical Engineering* 121(6):605-611, December 1999.

- S. E. Salcudean, personal communication, September 2001.
- B. M. Schena, personal communication, August 2001.
- C. L. Schuler, "Actuator Having Electronically Controllable Tactile Responsiveness," U.S. Patent No. 5,220,260, June 15, 1993.
- C. L. Schuler, "Actuator Having Electronically Controllable Tactile Responsiveness," U.S. Patent No. 5,414,337, May 9, 1995.
- C. L. Schuler, "Actuator Having Electronically Controllable Tactile Responsiveness," U.S. Patent No. 5,559,412, September 24, 1996.
- SolidWorks Corporation, <http://www.solidworks.com>, 300 Baker Avenue - Concord, MA 01742, August 2001.
- J. J. Slotine and W. Li, "Applied Nonlinear Control," Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1991.
- M. A. Srinivasan, R. J. Gulati, and K. Dandekar, In vivo compressibility of the human fingertip, *Advances in Bioengineering*, Ed. M.W. Bidez, ASME Winter Annual Meeting, November 1992.
- I. Sutherland, "The Ultimate Display," in *Proc. of the IFIP Congress*, pp. 506-508, 1965.
- The Mathworks, Inc., "Using Matlab: Version 6," Chapter 15: Differential Equations, product documentation, available from http://www.mathworks.com/access/helpdesk/help/pdf_doc/matlab/using_ml.pdf, The Mathworks, Inc., Natick, MA, 2001, p. 15-1.
- The Mathworks, Inc., "MATLAB Function Reference, Volume 2: F-O, Version 6," product documentation, available from http://www.mathworks.com/access/helpdesk/help/pdf_doc/matlab/refbook2.pdf, The Mathworks, Inc., Natick, MA, 2001, pp. 2-638 – 2-647.
- Trust Automation, San Luis Obispo, CA, personal communication, July 2001.
- J. C. Tsai and J. E. Colgate, "Stability of discrete time systems with unilateral nonlinearities," *Proceedings of the International Mechanical Engineering Congress and Exposition*, ASME, Chicago, IL, DSC Vol. 57-2, pp. 695-702, November, 1995.
- Y. Z. Tsytkin, "On the Stability in the Large of Nonlinear Sampled Data Systems," *Dokl. Akad. Nauk USSR*, Vol. 145, pp. 52-55.

- R. T. Verrillo, "A Duplex Mechanism of Mechanoreception," In D. R. Kenshalo, editor, *The Skin Senses*, pp. 139-156, Charles C. Thomas, Springfield, IL, 1968.
- F. E. Zajac and J. M. Winters, "Modeling Musculoskeletal Movement Systems: Joint and Body Segmental Dynamics, Musculoskeletal Actuation, and Neuromuscular Control," *Multiple Muscle Systems: Biomechanics and Movement Organization*, Chapter 8, J. M. Winters and S. L-Y. Woo (ed.), Springer-Verlag, New York, 1990.
- C. B. Zilles and J. K. Salisbury, "A Constraint-based God-object Method for Haptic Display," *Proc. IEEE/RSJ Int. Conf. On Intelligent Robots and Systems*, Pittsburgh, PA, 1995, pp. 146-151.