

APPLICATIONS OF FEEDBACK CONTROL
TO MUSICAL INSTRUMENT DESIGN

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

Edgar Joseph Berdahl
December 2009

© 2010 by Edgar Joseph Berdahl. All Rights Reserved.

Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial-No Derivative Works 3.0 United States License.
<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/pb603fs3989>

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.

Julius Smith, III, Primary 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.

Gunter Niemeyer

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 Salisbury, Jr

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumpert, Vice Provost Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Preface

In 1963, Max Mathews wrote that the “numbers which a modern digital computer generates can be directly converted to sound waves. The process is completely general, and any perceivable sound can be so produced” [106]. Since roughly this time, composers have been able to synthesize an amazing variety of sounds using diverse techniques [134]. However, compared to the development of traditional music over several centuries, the development of computer music has been so quick and so eager that the quality of the physical interaction between the performer and the instrument has deteriorated, both to the detriment of the performance and its reception by the audience.

By now laptops have become computationally powerful yet relatively inexpensive and portable. It is no surprise that many computer music performers experiment with the sound synthesis and sensing capabilities offered to them by laptops. Indeed, this approach is the easiest and most cost effective approach. Let us consider a typical live performance of modern computer music. A performer might cause a loudspeaker to emit a musical sound wave by alternatively pressing the keys, shaking the accelerometer, and/or sliding his or her finger across the trackpad of a laptop. Yet the relationship between these small movements and the sound may not be immediately obvious. As a consequence, the members of the audience or other performers on stage may not build an accurate conceptual model describing the musical interaction. What do they perceive about the quality of the performance? How can they judge the degree of virtuosity of the performer? How quickly is the performer able to improve his or her own playing level? How can we make it easier for people to build a conceptual model of a novel musical interaction?

I argue that we need to search for satisfying answers to these big questions now that we have harnessed the power of computation and sensing. Any arbitrary programmed interaction involving sensors, computers, and loudspeakers will not necessarily be immediately comprehensible because there is an overabundance of design freedom. Rather than being haunted by this overabundance, known by some as “programmability is a curse” [53], we should develop designs that promote the performers’ and audience’s forming of accurate conceptual models that describe the musical interaction.

At present, I cannot provide general answers to all of the questions posed above. However, I would like to point out that Physical Interaction Design presents us with a design method that addresses these questions adequately. Physical Interaction Design is predicated on the fact that humans are physical beings, and we are accustomed to interacting with objects in the physical environment [165][111]. We frequently build and refine internal conceptual models for these interactions [169], and we can construct higher order models out of more basic models that we have already developed [139]. We can use these models to transfer skills between virtual and real environments [86]. When we design musical instruments that incorporate meaningful physical interactions, whether these interactions are implemented in the real world or in a virtual world, these instruments promote the formation of rich conceptual models.

Some computer music performance paradigms are beginning to reflect Physical Interaction Design and in particular Newton’s third law, which states that “for every action there is an equal and opposite reaction.” We can re-interpret Newton’s third law rather loosely and abstractly as follows: when a performer manipulates a sensor to cause a sound to be synthesized, the resulting sound wave should emanate from an actuator near the sensor. For example, laptop orchestras are already widely adopting the concept that each musician’s sound source should be collocated with the musician’s physical actions [159]. This approach requires more cables and amplifiers; however, it makes it easier for performers and audience members alike to understand where sounds are coming from and why they are being synthesized. As a consequence, this approach helps the performers and the audience form accurate

conceptual models.

I did not plan for the instrument designs outlined in this thesis to conform to the notions of Physical Interaction Design. I merely set out with the idea of enhancing the musician’s interaction with the instrument by providing improved force feedback. In other words, my goal was to study how to make electronic musical instruments more tangible, so that the performance of electronic music would seem less artificial, “dry,” “synthetic,” and “mechanical” [29]. Starting out from an engineer’s point of view, I studied classes of controllers for generating force feedback. Positive real controllers were guaranteed to result in a stable control system when connected to any dissipative acoustic musical instrument. To my great surprise, I discovered that any positive real controller is equivalent to a passive mechanical system. In other words, a large set of controllers that are guaranteed to be stable when applied to any dissipative acoustic musical instrument, correspond to physical analog models!! From this point onward, I repeatedly found that physical analog systems provided solutions to the practical control problems I studied. As a consequence, I have become an unexpected fervent proponent of Physical Interaction Design.

I would like to push the idea of applying Physical Interaction Design to Newton’s third law a bit further. Let’s consider a more strict interpretation of Newton’s third law. We assume now that a sensor and an actuator are perfectly collocated in the mathematical sense described in Section 2.3.2. Now when the performer pushes on the sensor, the collocated actuator pushes back on the musician. In other words, while the performer is in contact with the musical sounding object, he or she feels the natural, structural vibrations of the object that he or she is inducing. The physical interaction is collocated with the result—the sound wave emanates from the position of the collocated sensor and actuator. I argue that *sound and touch should be collocated*.

I discuss two paradigms for designing these types of musical sounding objects. In chapters 1-4, I explain feedback controlled acoustic musical instruments, which enable a musician to interact with a traditional, acoustic musical instrument, whose acoustical properties are modified by collocated feedback. The performer interacts directly with the physical medium responsible for sound production.

Haptic musical instruments are presented in chapters 5-7. A large class of collocated haptic musical instruments enable the performer to interact with virtual musical instruments. These instruments promote natural physical interactions and motivate the guidelines I provide in Chapter 8.

Dedication

This is dedicated to my parents Paul Berdahl and Jane Berdahl and my brother Carl Berdahl for their love and support.

Acknowledgements

I would like to extend my most heart-felt thanks to my first two thesis readers Julius O. Smith III (Music/EE) and Günter Niemeyer (ME/EE) for their ongoing encouragement and support. I came to Stanford University specifically to work with Julius, and he never ceased to continue inspiring me with his seemingly never-ending knowledge about digital signal processing and sound synthesis. Julius gave me the freedom to study unconventional yet fundamental approaches to synthesizing musical sound. He supplied me with financial support from CCRMA and the RealSimple project, and when I needed physical laboratory space, he generously provided me with access to one of the desks in his own office.

Without Günter's guidance, this thesis would not have been possible either. First, he aided me by opening up my eyes to the practical knowledge necessary to make feedback control systems work well in practice. In addition, he taught me a great deal about how to organize and develop research in a rigorous framework.

I would further like to thank Kenneth Salisbury (CS/Surgery), my third reader, especially for his emphatic support in studying new applications of haptics to synthesizing sound. Finally, I thank Scott Klemmer, my committee chair, for encouraging me to think about how my thesis work relates to the principles of human-computer interaction.

Many other researchers have influenced me during my dissertation research. Most notably, Bill Verplank and Chris Chafe inspired me to consider how to make it easier to play physical models in real performance contexts. Bill and I had many valuable discussions about prior work and new ideas, and Chris provided me with the means to get involved with co-teaching the “Physical Interaction Design for Music” course

sequence at CCRMA. I would like to further thank: Wendy Ju, who co-taught with me at CCRMA; Jonathan Abel, who continues to serve as a great inspiration in the field of acoustics and signal processing; Veronique Larcher, who made it possible for me to study at Sennheiser Research; John Chowning, who bless his soul always knows just what to say; David Wessel, who also always knows just what to say and hosted me at CNMAT; Adrian Freed, who helped me track down prior work on my thesis topic in France; Charles Besnainou, a luthier who made fundamental research in sensors and actuators for feedback controlled acoustic musical instruments; Tom Rossing, who coauthored the book I always cite on the acoustics of traditional musical instruments; Marty Johnson, Dave Berners, Max Matthews, Jonathan Berger, Michael Gurevich, Joan Savall, Luiz Souza, Kyle Machulis, Heiko Zeuner, and Peter Lindener, may his infinite enthusiasm never wane for changing the world with feedback control.

The CCRMA staff have been a great help: Tricia Schroeter, Carr Wilkerson, Fernando Lopez-Lezcano, Sasha Leitman, Jay Kadis, Chryssie Nanou, Linnea Williams, and Sandy Greenfield. Besides the energetic master's degree students passing through CCRMA, I would like to thank the longer term graduate students: David Yeh, Nelson Lee, Tim Stilson, Matt Wright, Ryan Cassidy, Kyogu Lee, Hiroko Terasawa, Harvey Thornburg, Gautham Mysore, Steven Backer, Per Bloland, Rob Hamilton, Bruno Ruviano, Juan-Pablo Caceres, Juhan Nam, Greg Sell, Aaron Master, (Tak) Pamornpol Jinachitra, Woon Seung Yeo, Sook Young Won, Luke Dahl, Nicholas Bryan, Song Hui Chon, and Hoi Wong. I would also like to thank students in Günter Niemeyer's lab including Probal Mitra and June Park, the Stanford Graduate Fellowship program, and anyone that I may have forgotten.

Contents

Preface	v
Acknowledgements	x
1 Introduction	1
1.1 Musical Instruments	1
1.2 Importance Of Haptic Feedback	3
1.3 Feedback Controlled Acoustic Musical Instruments	4
1.4 Haptic Musical Instruments	4
1.5 Contributions	7
1.5.1 Feedback Controlled Acoustic Musical Instruments	7
1.5.2 Haptic Musical Instruments	11
2 Feedback Controlled Instruments	14
2.1 Outline	14
2.1.1 Overview	14
2.1.2 Musical Acoustics Review	15
2.1.3 Problem Description	16
2.1.4 Challenges	18
2.1.5 Prior Examples	19
2.2 Basic Models	20
2.2.1 Single Resonance Model Of Instrument	20
2.2.2 Waveguide Model	21
2.3 Passive Control	22

2.3.1	Definition	22
2.3.2	Positive Real Functions	23
2.4	Control System Stability	24
2.4.1	Revised Bode Stability Criterion	25
2.4.2	Stability Proof	25
2.4.3	Implications	26
2.5	Practical Collocation Considerations	27
2.5.1	Matched Sensor and Actuator	27
2.5.2	Sensoriactuators	29
2.5.3	Digital Control System Delay	30
2.6	Organization	30
3	Changing The Linear Dynamical Properties	34
3.1	Passive Control	34
3.1.1	Proportional-Integral-Derivative (PID) Control	37
3.1.2	Changing The Resonance Frequency Is Difficult	39
3.1.3	Single Resonance Model Predicts General Behavior	43
3.1.4	Controller Phase Response	51
3.1.5	Bandpass Control	51
3.1.6	Notch Filter Control	52
3.1.7	Alternating Filter	55
3.1.8	Teleoperator-Based Control	63
3.1.9	Positive Real Multiple-Input Multiple-Output Controllers	64
3.2	Positive Real Controllers Motivate The Design Of Nonpassive Controllers	72
3.2.1	Negative Spring Constant	72
3.2.2	Negative Damping	72
3.2.3	Negative Mass	73
3.2.4	Integral Control	73
3.3	Higher Order Control Schemes	75
3.3.1	Non-Collocated Traveling Wave-Based Control	75
3.3.2	Termination-Based Control	83

3.3.3	Loop Gain Limitations In The Laboratory	91
3.3.4	Matched Termination Control	97
3.4	Summary	100
4	Inducing Self-Sustaining Vibrations	103
4.1	Introduction	103
4.1.1	Self-Sustaining Oscillators	104
4.1.2	Sources of strong nonlinearity in acoustic musical instruments are often localized	104
4.2	Application To Musical Self-Sustaining Oscillators	105
4.3	Event-Based Control	109
4.3.1	Approach	109
4.3.2	Bandwidth Advantage	112
4.3.3	Analysis	112
4.3.4	Event Detector	115
4.4	RMS Level-Tracking Controllers	117
4.4.1	Dynamic Range Limiter Control	117
4.4.2	RMS Level-Tracking Controllers	119
4.4.3	Adaptive RMS Level-Tracking Controllers	127
4.5	Summary	134
4.6	Conclusions	134
4.7	Implementation	136
5	Haptic Musical Instruments	138
5.1	Introduction	138
5.2	Overview	138
5.3	Controlling Acoustics	140
5.4	Controlling Gestures	141
5.4.1	Motivation	141
5.4.2	Who Is In Control?	144
5.4.3	Philosophy	145
5.4.4	Human Motor Control	147

5.4.5	Closed-Loop Versus Open-Loop Control	148
5.4.6	Haptic Assistance In General	150
5.4.7	Haptic Assistance For Playing Rhythm And Pitch	150
5.5	Overview	151
6	Making New Gestures	153
6.1	Simple Model	154
6.1.1	Above The Drum Membrane	154
6.1.2	Simple Collision Model	155
6.1.3	Upward Soft Collisions	156
6.2	Haptic Drumstick	157
6.2.1	Overview	157
6.2.2	Introduction	157
6.2.3	Altering The Physical Interaction	159
6.2.4	Stability Analysis	160
6.2.5	Results	162
6.2.6	Limitations	163
6.3	Haptic Drum	163
6.3.1	Prototype v1 Design	164
6.3.2	Dynamic Behavior	166
6.3.3	Prototype v2 Design	169
6.4	Conclusions	171
7	Making More Accurate Gestures	175
7.1	More Accurate Rhythmic Elements	175
7.2	More Accurate Pitches	176
7.3	Pitch Selection Problem Overview	177
7.4	Prior Work	178
7.5	Simple Theremin-Like Instrument	179
7.6	Haptic Assistance	180
7.6.1	Basic Detent	180
7.6.2	Force Feedback Conditions	181

7.6.3	Further Ideas From Literature	185
7.7	Subject Test	186
7.7.1	Experiment Design	186
7.7.2	Data Analysis	189
7.7.3	Selecting The Parameters For DET	193
7.7.4	Qualitative Results	198
7.7.5	Results	201
7.8	Final Words	207
8	Conclusions	209
8.1	Review	209
8.2	Collocation	209
8.3	Physical Analog Models	210
8.4	Further Advantages of Active Feedback	210
8.5	Access to Feedback Technology	211
8.6	Haptic Assistance Guidelines	211
A	Haptic Digital Waveguide String	213
A.1	Introduction	213
A.1.1	Prior Work	214
A.1.2	Finite Difference Modeling	214
A.1.3	Finite Element Modeling	214
A.1.4	Digital Waveguide Modeling	215
A.2	Direct Modeling Approach	215
A.2.1	Explicit Control Junction Element	216
A.2.2	Interfacing With The Waveguide	217
A.2.3	Implicit Control Junction Element	219
A.3	Final Words	220
B	Further Information On Subject Test	223
B.1	Documents	223
B.2	Questionnaire Comments	223

C	HSP Toolbox: An Overview	228
C.1	Introduction	228
C.2	Prior Work	229
C.3	Hardware	230
C.4	1DOF Examples	231
C.4.1	Basic Elements	231
C.4.2	Haptic Drum	232
C.4.3	Detents and Force Profiles	232
C.5	Multi-DOF Examples	233
C.5.1	Surface	233
C.5.2	Models	234
C.6	Performing Telemusic	236
C.7	Conclusions	237
D	Nonlinear PID Control	238
D.1	Overview	238
D.2	Damper	239
D.3	Spring	240
D.4	Lyapunov Stability Of Origin	241
D.5	Implications Of Passivity	242
E	Additional Musical Instruments	243
E.1	Tangible Virtual Vibrating String	243
E.1.1	A Physically Motivated Virtual Musical Instrument Interface .	243
E.1.2	Introduction	243
E.1.3	Prior Guitar-Like Interfaces	245
E.1.4	Tangible Guitar String	246
E.1.5	Website	250
E.1.6	Future Work	250
E.1.7	Conclusion	251
E.2	Electromagnetically Prepared Piano	251

F	Actuator and Sensor Design	258
F.1	Harmonic Distortion	258
F.2	Actuators	259
F.2.1	Lorentz' Law Type	259
F.2.2	Reluctance Force Type	261
F.3	Sensors	266
F.3.1	Optical Displacement Sensors	266
F.3.2	Piezoelectric Sensors	273
F.3.3	Sensor THD Estimation	274
F.4	Collocation Considerations	276
F.5	Active String Termination Concept	277
G	Toolbox for the Feedback Control of Sound	279
G.1	Introduction	279
G.2	Open Source Solution	281
G.3	Operation	282
G.4	Estimated PDFs Revealing Jitter Performance	286
G.5	Conclusion	289
G.6	Initialization Files	290
G.6.1	grub.conf	290
G.6.2	/usr/realtime/init_rtai	290
G.6.3	/usr/realtime/start_rtai	291
G.6.4	/usr/local/comedi/load_cmodes	291
G.6.5	kcomedilib.c missing functionality	292
	Bibliography	293

List of Tables

2.1	Classification of acoustic musical instrument types	31
3.1	Relation between distortion and maximum safe waveguide loop gain .	82
5.1	System delay times for the most important active feedback mechanisms of human motor system	148

List of Figures

1.1	Most basic human-musical instrument interaction model	2
1.2	Human musical-instrument interaction model with all feedback loops completed	2
1.3	Musician interacting with a feedback controlled acoustic musical instrument	5
1.4	Musician interacting with a typical haptic musical instrument	5
1.5	Musician interacting with an idealized haptic musical instrument (see footnote 1)	6
2.1	Musician interacting with a feedback controlled acoustic musical instrument	17
2.2	Simple block diagram for feedback control of an acoustic musical instrument	18
2.3	Damped harmonic oscillator with mass m , damping parameter R , spring constant k , and external force $F(t)$	20
2.4	Digital waveguide model of a vibrating string with one sensor and one actuator	21
2.5	Relative dimensions of string model	22
2.6	Driving point of a system	23
2.7	Simple block diagram of control system with $s = j\omega$	26
2.8	Overview of controller classes covered in this document	33
3.1	Collocated control of a vibrating string with a single sensor and actuator	35
3.2	Summary of positive real controllers covered in Section 3.1	36

3.3	Collocated control of a vibrating string with a collocated sensor (depicted as microphone) and and actuator (depicted as loudspeaker) (top); Mechanical analog of PID control (bottom)	39
3.4	Root locus for $P_P > 0$ proportional to displacement control on left and root locus for $P_D > 0$ derivative of displacement control on right	46
3.5	Root locus for $P_{DD} > 0$ second derivative of displacement control on left and root locus for $P_I > 0$ integral of displacement control on right	48
3.6	Frequency response of the bandpass controller given by (3.31) with $Q = 20$ and $w_c = 2\pi 200\text{Hz}$	53
3.7	Mechanical analog of the bandpass controller given by (3.31)	53
3.8	Frequency response of the notch filter controller given by (3.32) with $w_c = 2\pi(300\text{Hz})$, $Q = 1$, and $\alpha = 10$	54
3.9	Mechanical analog of the notch controller given by (3.32) (left); Pole-zero plot of (3.32) (right)	55
3.10	Frequency response of $K_{af}(z)$ where the x 's mark where the transfer function overlaps with the resonances of the waveguide instrument with fundamental frequency $f_0 = 100\text{Hz}$	56
3.11	Pole zero plot of the digital implementation of $K_{af}(z)$ (zoomed in) . .	57
3.12	Pole zero plot of the digital implementation of $K_{af}(z)$	58
3.13	Frequency response of the alternating filter $K_{af}(z)$	60
3.14	Frequency response $K_{af}(z)z^{-1}$ incorporating the controller delay . . .	61
3.15	Fourier transform of windowed open loop response (solid blue) and closed loop response (dash-dotted red)	62
3.16	Teleoperator-based control of a vibrating string	63
3.17	Mechanical analog for the controller described by (3.33)	65
3.18	The gyrator element	66
3.19	Block diagram for gyrator connecting two independent second-order systems	67
3.20	Root locus for gyrator applied to two independent lightly-damped harmonic oscillators	68
3.21	Gyrator connected between the points A and B along a vibrating string	69

3.22	Mobility at junction B of the system shown in Figure 3.21	71
3.23	Root locus over $P_I < 0$ for $K_{int}(s)$ applied to a lightly damped harmonic oscillator	74
3.24	Representation of a physical waveguide	76
3.25	Configuration for non-collocated traveling wave-based control	77
3.26	Configuration after physical reflection has already been damped.	78
3.27	Configuration where sensor nonlinearity is modeled by $\frac{\partial \dot{v}_R}{\partial v_R}$	79
3.28	Configuration for controlling the termination's transfer reflection function	85
3.29	$ K_{inverting}(f) $ for $\tau_A = 0$ (collocated, thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)	88
3.30	$ K_{inverting,f_0}(f) $ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)	89
3.31	$ K_{inv,practical}(f) $ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)	89
3.32	$ K_{noninvert}(f) $ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue) .	90
3.33	Dashpot attached to a rigidly-terminated semi-infinite vibrating string	92
3.34	Equivalent mechanical configuration to Figure 3.33 where the impedance of the right-hand string segment is lumped into the dashpot	92
3.35	Dashpot attached to a rigidly-terminated vibrating string with reflection coefficients labeled for damping analysis	94
3.36	Increase in loop gain needed for changing the pitch rather than the damping as a function of frequency f_0 ($\tau_{damped} = 15$ ms and $\tau_{change f_0} = 120$ ms)	95
3.37	Active termination control using a perfectly-terminated string	98
3.38	Active termination control using a piston-type collocated sensor and actuator	99
4.1	Nonlinear damping control of a vibrating string (top), and string coupled to a nonlinear damper (bottom)	106
4.2	Damping characteristic curve due to holding a stationary bow held against a string	107
4.3	Bowed string nonlinear damping characteristic	107

4.4	Clarinet reed nonlinear damping characteristic (left) and trumpet lip nonlinear damping characteristic (right)	108
4.5	Configuration for event-based control of a vibrating string	109
4.6	Example of traveling velocity waves on a string actuated with a short pulse	111
4.7	Time domain (left) and frequency domain (right) representation of the p , v , and m signals for event-based self-sustaining control of a string. (Note that the illustrated $P(f)$ is not the exact Fourier transform of $p(t)$.)	114
4.8	Control configuration with dynamic range limiter in the feedback loop	120
4.9	RMS level-tracking controller	121
4.10	RMS level-tracking controller (RMS signals only)	122
4.11	Valid region for convergent controllers (left); Valid region for exponentially- convergent controllers (right)	124
4.12	Trajectories for the dynamic range limiter controller (left); Trajectories for the controller supporting multiple equilibria (right)	125
4.13	Adaptive RMS level-tracking controller	128
4.14	Ten plucks of linearly increasing magnitude (above), corresponding loop gain $L(t)$ (below)	130
4.15	Driving the RMS of the second harmonic to two values	131
4.16	System block diagram for implementing multiple RMS level-tracking controllers concurrently	132
4.17	RMS levels for resonant ring modulation	134
4.18	Summary of nonlinear controllers covered in Chapter 4	135
4.19	Feedback Guitar	137
4.20	Robert Hamilton playing the guitar	137
5.1	Musician interacting with a typical haptic musical instrument	139
5.2	Paradigms for musical practice	142

5.3	Two extremes: a robot playing a musical instrument (left) and a human playing a musical instrument (right); and the spectrum in between (arrow)	144
5.4	A musician is assisted by haptic devices	147
6.1	Drumstick dynamics for drumstick above membrane	155
6.2	Drumstick dynamics for $z < 0$	156
6.3	Graphical display (above) and PHANTOM robotic arm (lower right)	158
6.4	Block diagram describing the operation of the haptic drumstick . . .	158
6.5	Example system trajectories in the phase plane	161
6.6	Haptic Drum v1	164
6.7	Haptic Drum prototype v1 with additional percussion instruments as used for performing “Percussion Edstravaganza”	165
6.8	Vertical position of the drumstick while a musician plays a regular drum roll at 70Hz	167
6.9	Vertical position of the drumstick shown for complex behavior	168
6.10	Drumstick position (top) and velocity (bottom) for a low-rate drum roll	169
6.11	Example phase plane trajectory for two impacts corresponding to Figure 6.10	170
6.12	Drumstick position (top) and velocity (bottom) for a single-handed haptically assisted double stroke roll	170
6.13	Example phase plane trajectory for four impacts corresponding to the data shown in Figure 6.12	171
6.14	Haptic Drum prototype v2 with built-in synthesizer	172
6.15	Haptic Drum prototype v2 used with standard drum set	172
6.16	Promotional flyer for the Haptic Drum v2	173
7.1	Drumstick position (above) and velocity (below) for a regular drum roll	176
7.2	Musician’s finger inserted into the PHANTOM thimble	180
7.3	Force profile $F(y)$ and terrain height profile $h(y)$	181
7.4	Force profile for the DET condition	182
7.5	Strength of haptic assistance for the VEL condition	183

7.6	Strength of haptic detents for the FRC condition	184
7.7	Melody excerpts	188
7.8	Fundamental frequency contour for subject 7 and melody excerpt 5 for the NOFB force condition (whole unsegmented contour in thin black line, segmented contour of notes played in thick blue line, and estimated intended notes shown in thick dash-dotted red line)	190
7.9	RMS acceleration of fundamental frequency contours for subject 7 and melody excerpt 5 for the NOFB force condition (TOP: unsegmented fundamental frequency contour in thin, dash-dotted black line; seg- mented contour in thick, blue line; BOTTOM: unsegmented accelera- tion of fundamental frequency contour in thin, dash-dotted black line; segmented contour in red)	192
7.10	Basic detent element	193
7.11	Fundamental frequency contour played by one of the authors subject to the DET condition for $D = 0.5\text{N}$ and $\rho = 0.3$	194
7.12	Sweep 1: MAE for one of the authors playing the diatonic scale with $\rho = 0.3$ while D was swept from about zero to about 1.2N	195
7.13	Sweep 1: RMS acceleration for one of the authors playing a glissando with $\rho = 0.3$ while D was swept from about zero to about 1.2N	196
7.14	Sweep 2: MAE for one of the authors playing the diatonic scale with $D = 0.41\text{N}$ while ρ was swept from nearly zero to 0.5	196
7.15	Sweep 2: RMS acceleration for one of the authors playing a glissando with $D = 0.41\text{N}$ while ρ was swept from nearly zero to 0.5	197
7.16	Finalized basic detent element shape with $\rho = 0.2$, $W = 0.375\text{cm}$, and $D=0.41\text{N}$	198
7.17	NOFB No Haptic Feedback (subject 7, melody 5, $MAE=0.43$, $Accel=9$)	199
7.18	SPR Spring Feedback (subject 7, melody 5, $MAE=0.43$, $Accel=14$) .	200
7.19	DET Multiple detents (subject 7, melody 5, $MAE=0.23$, $Accel=16$) .	200
7.20	FRC Force-Sensitive Detents (subject 7, melody 5, $MAE=0.17$, $Accel=10$)	201
7.21	Box plot of the mean absolute error (MAE) of note fundamental fre- quency trajectories across all subjects and melody excerpts	202

7.22	Box plot of the RMS acceleration of the fundamental frequency trajectories of glissandi across all subjects and melody excerpts	203
7.23	Preferences of the test subjects	204
7.24	Fundamental Frequency MAE shown as a function of time during the test trials to measure residual learning during testing	206
7.25	Box plot of fundamental frequency MAE across the test melody excerpts	207
A.1	Signal flow diagram for hybrid modeling approach	216
A.2	Signal flow diagram for direct modeling approach	216
A.3	Explicit control junction connected to a digital waveguide model of a vibrating string	217
A.4	Teleoperator control of a virtual vibrating string	218
A.5	Signal flow diagram detailing how to implement teleoperator control .	218
A.6	Helmholtz motion of haptically bowed virtual string	219
A.7	Implicit control junction connected to a digital waveguide model of a vibrating string	221
A.8	Plucking and damping using the direct approach	222
B.1	Explanation provided to test subjects prior to the procedure	224
B.2	Questionnaire	225
C.1	Musician interacting with a haptic musical instrument	229
C.2	Musician holding the Falcon grip	230
C.3	Spring implementation in HSP	231
C.4	Force profile $F(y)$ (above) and virtual terrain height profile $h(y)$ (below) for a simple detent.	233
C.5	A force profile stored in a graphical <i>array</i> in Pd (stable positions circled in blue)	234
C.6	A surface modeled with a user-editable graphical array object in Pd .	235
C.7	Falcon with shaker attached to grip	236
D.1	Linear damping (left) and saturating damping (right)	239

D.2	Damping characteristic curve due to holding a stationary bow held against a string	240
D.3	Stiffening spring characteristic curve	241
E.1	Tangible virtual string signal flow diagram	246
E.2	Two string segment approach	248
E.3	Two axis digital waveguide string model	249
E.4	Graphtech piezoelectric pickup	249
E.5	Tangible Guitar String Interface	254
E.6	Close up of the motors	255
E.7	System block diagram for one channel	256
E.8	Electromagnetically prepared piano prototype	257
F.1	SISO system	259
F.2	Lorentz' law type actuator	260
F.3	Lorentz' law type actuator (viewed from above)	261
F.4	Lorentz' law type actuator (oblique angle)	262
F.5	Hand-wound reluctance force actuator	262
F.6	Relationships between f_X , i , and x	263
F.7	Expected magnetic field lines for a different geometry (wider core, multiple permanent magnets)	264
F.8	Photodiode type optical sensor in the laboratory (LED not shown)	267
F.9	Photodiode type optical displacement sensor	267
F.10	Halogen lamp	269
F.11	Nonuniform radiation pattern of a coiled halogen lamp	270
F.12	More uniform radiation pattern of lamp	271
F.13	Piezoelectric guitar string termination sensor from Graphtech	272
F.14	Piezoelectric guitar string termination sensor from Graphtech	274
F.15	Spectrum analyzer showing the nonlinearity in the sensor signal for a S2387-1010R silicon photodiode illuminated by the halogen lamp	275
F.16	Collocation geometry	277
F.17	String termination actuator concept	278

G.1	Feedback control of a vibrating string (controller model includes a delay of D seconds)	279
G.2	Modular design of TFCS's software components	283
G.3	The GUI for TFCS provides buttons, toggles, sliders, etc.	284
G.4	The GUI also allows arbitrary nonlinear characteristic curves for controller elements to be drawn with the mouse.	285
G.5	A continuous time signal (in blue) as approximated by a zero-order hold (dash-dotted in red)	285
G.6	Estimated jitter for TFCS given no additional load ($f_S = 40\text{kHz}$) . .	287
G.7	Estimated jitter for TFCS given significant system load ($f_S = 40\text{kHz}$)	288
G.8	Estimated jitter for TFCS given a system loaded by constantly dragging a window back and forth ($f_S = 40\text{kHz}$, 0.8% overruns)	288
G.9	Estimated jitter for TFCS with no load ($f_S = 50\text{kHz}$, 4.6% overruns)	289

Chapter 1

Introduction

1.1 Musical Instruments

Modern musical instrument design stems from a rich history, which is too long and varied to summarize here. Suffice it to say that traditional musical instruments were designed to provide a human with intimate control over mechanical elements used to induce audible sound vibrations [85][62]. In this thesis, we refer to these instruments as *acoustic musical instruments*. In contrast, modern technological developments in sensing, electrical engineering, and computer engineering have freed musical instrument designers from traditional limitations, paving way for the development of radically different, new musical instruments. Because these instruments are developing so quickly and provide so much design freedom [53], it is important that we step back from the musical instruments that are possible and focus on those instruments that promote a rich interaction between the musician and the instrument.

The most basic interaction model for a musician playing a musical instrument is often the following: a musician presents a musical instrument with a mechanical excitation while receiving a response in the form of auditory feedback. This interaction model is shown in Figure 1.1, where we assume for convenience that the musician uses his or her hands to provide the mechanical excitation. We acknowledge that a wide range of other excitations can of course be provided using the mouth, the remainder of the body, etc. Figure 1.1 is drawn according to Bill Verplank’s physical interaction




Figure 1.1: Most basic human-musical instrument interaction model




Figure 1.2: Human musical-instrument interaction model with all feedback loops completed

design schema [165][111], except that the labels emphasize physical parameters. Other researchers in the field of computer music employ diagrams that emphasize the separation of the musical instrument interface from the sound synthesizer [110][164][40]; however, we prefer our more general depiction because it emphasizes the musical interaction. It is particularly important to draw the musician's hand [46].

To make the depiction of the musical interaction more precise, we need to include additional feedback channels that the musician uses. For example, besides auditory feedback, a musician uses visual and haptic feedback, for instance to orient his or her hands on a piano keyboard. A musician also uses his or her brain to control the mechanical excitation of the instrument. Accordingly, Figure 1.2 includes the brain as the closing link in the composite feedback system, which includes visual, auditory, and haptic feedback. The internal loop between the brain and the hand illustrates the brain governing the human motor control system [139].

1.2 Importance Of Haptic Feedback

There is a general consensus within the research community that auditory and visual feedback are helpful for a musician performing in a live scenario. Auditory feedback is necessary for the musician to ensure that the musical outcome is what he or she expects. Visual feedback is widely known to be useful for completing tasks in general. However, the importance of haptic, tactile, vibrotactile, and mechanical feedback is sometimes underestimated, and it tends to be neglected in the design of new musical instruments.

Many traditional musical instruments provide musicians with haptic feedback. For instance, most traditional acoustic instruments vibrate significantly, allowing the musician to receive information relating to the state of the instrument. Even instruments such as the pipe organ, where the musician is not directly coupled to the mechanically vibrating elements, have keys that the musician touches and feels while playing. The Theremin is a notable exception. It is played by a musician moving his or her hands about in free space, hence it does not really provide any significant mechanical or haptic feedback [156]. The Theremin is also known to be difficult to

play accurately [120].

Besides work in this thesis (see Chapters 5-7), other studies have indicated that haptic feedback for musical instruments is vital. For instance, studies of reaction times have indicated that humans can react faster to haptic stimuli than to auditory and visual stimuli [128]. Some studies have focused specifically on musical instrument performance. In 1993, Chris Chafe attached an accelerometer to the fingernail of a cellist’s left index finger. The finger was used to terminate the string and measurements were made suggesting that the cellist player should be able to distinguish various vibrotactile cues describing the motion of the bowed string [47]. In other words, a cellist should be able to use haptic feedback to distinguish string transients from other string states. Birnbaum completed an even more extensive study describing the roll of vibrotactile feedback in musical instruments [30]. In 2000, O’Modhrain published results from an experiment suggesting that haptic feedback enables musicians to perform more accurately compared to a condition with no haptic feedback at all [120].

1.3 Feedback Controlled Acoustic Musical Instruments

In this thesis, we study how the musical interaction can be altered by adding new feedback loops to the diagram shown in Figure 1.2. For instance, chapters 2 through 4 detail how an additional feedback controller can be added to fundamentally change the acoustics of the musical instrument. Figure 1.3 illustrates the musical interaction promoted by feedback control of an acoustic musical instrument.

1.4 Haptic Musical Instruments

The final portion of the thesis relates to “haptic” musical instruments, which are introduced in Chapter 5. The word *haptic* is derived from the Greek word *haphe* meaning “pertaining to the sense of touch” [30]. A haptic musical instrument (HMI)

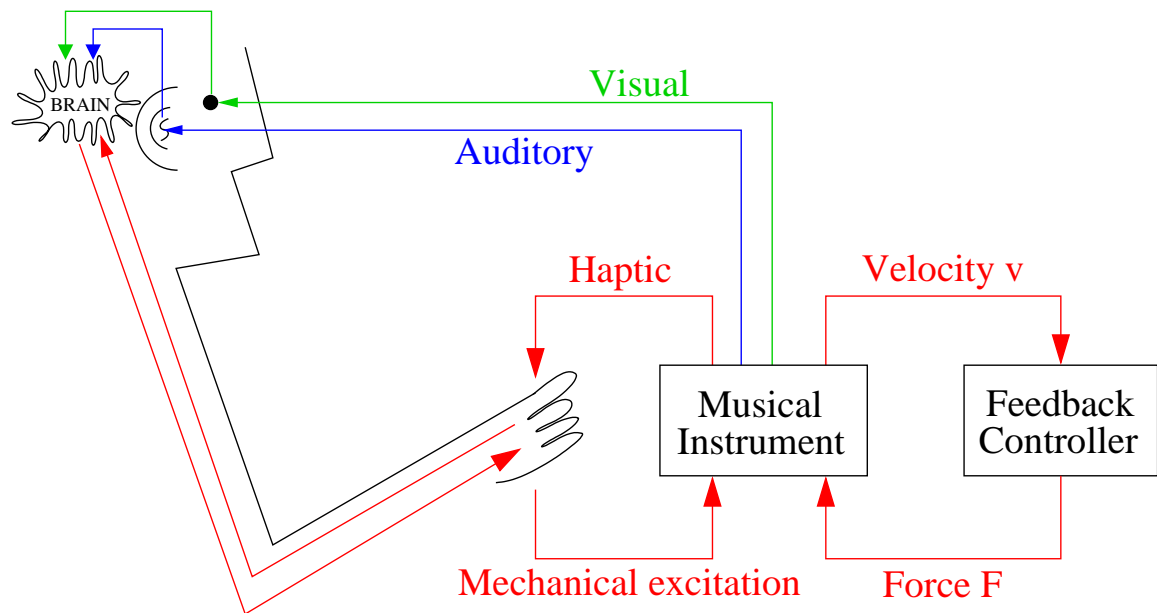


Figure 1.3: Musician interacting with a feedback controlled acoustic musical instrument

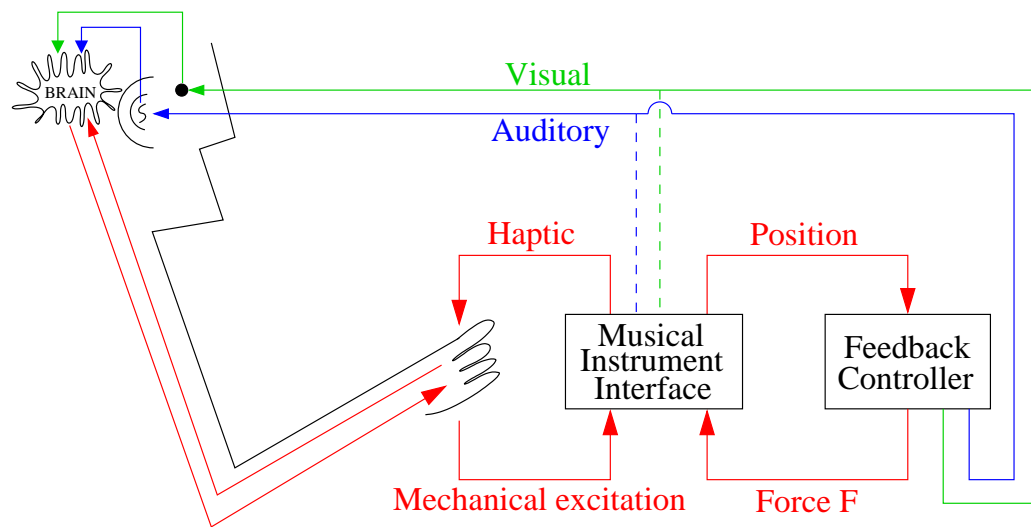


Figure 1.4: Musician interacting with a typical haptic musical instrument




Figure 1.5: Musician interacting with an idealized haptic musical instrument (see footnote 1)

is an electronic musical instrument that provides the musician not only with auditory feedback but also with haptic (or force) feedback. Figure 1.4 shows a representation of a musician interacting with a typical haptic musical instrument. The instrument consists of a musical instrument interface, which could for example be a robotic arm, whose dynamics are specified by a feedback controller.¹ In contrast with Figure 1.3, the primary forms of visual and auditory feedback are provided by the feedback controller, generally by way of a graphical display and a loudspeaker. The connection from the musical instrument interface to the visual and auditory feedback paths are shown dashed (see Figure 1.4) since depending on the application, they may play a less important role in affecting a musician's interaction with a haptic musical instrument.

Ideally the dynamics of the musical instrument interface are negligible (also known as “transparent”) compared to perceptual capabilities when the power for the musical instrument is turned off. In this case, the feedback controller can have a much

¹ In this example instrument, the feedback controller senses position and actuates force. However, alternate configurations exist where the controller would sense velocity, acceleration, or some other function of position as an input variable. In addition, admittance controllers are also possible, which often take *force* as an input variable and output velocity [2].

greater effect on the dynamics of the musical instrument interface when the power is turned on. In fact, in many cases, the feedback controller has so much authority over the dynamics of the musical instrument interface that we can imagine that the musician is directly connected to the feedback controller as depicted in the idealized representation shown in Figure 1.5. In other words, ideally the musician is directly coupled to a feedback controller that implements some dynamical system, which can be employed to implement a virtual reality system.

1.5 Contributions

This thesis contributes in two major areas.

1.5.1 Feedback Controlled Acoustic Musical Instruments

Positive Real Controller Design Framework (Section 2.4.3)

The dynamic behavior of acoustic musical instruments can be modified by feedback control. Many different sensor, actuator, and controller designs are possible. We argue that positive real controllers are fundamentally important because every mechanical system corresponds to a positive real controller. Consequently, in a control system with ideal sensors, actuators, and positive real controller, the loop gain can be made arbitrarily large without destabilizing the control system, as long as the acoustic musical instrument is dissipative. This implies that in practice, loop gains can be made relatively large, resulting in effective control of the acoustic musical instrument. In addition, the positive real controller design philosophy places requirements on sensors and actuators that greatly aid in designing viable hardware. In particular, each actuator should correspond to a collocated and matched sensor.

New Positive Real Controllers

We employ the positive real framework to design new positive real controllers for changing the dynamics of acoustic musical instruments. Where convenient, we also introduce the mechanical analogs of these controllers. To the field of the feedback

control of acoustic musical instruments, we introduce bandpass control (Section 3.1.5), notch filter control (Section 3.1.6), control using combinations of bandpass and notch filters, alternating filter control (Section 3.1.7), teleoperator-based control (Section 3.1.8), as well as some multiple-input multiple-output controllers including the gyrator (Section 3.1.9).

Integral Control Of The Displacement (Section 3.2.4)

Collocated Proportional-Integral-Derivative (PID) control of the velocity [28], which is positive real, suggests that feeding back signals that are proportional to the displacement, velocity, and acceleration can be useful. We introduce the concept of feeding back the integral of the displacement, which can be employed to change the decay times primarily of resonances at relatively low frequencies. It should be noted that the integration of the displacement must be computed in a “leaky” fashion in order to prevent destabilizing the acoustic musical instrument at 0Hz.

Generalized Predictions Of The Single Resonance Model (Section 3.1.3)

We introduce the term *musical resonance*, which is a resonance whose decay time is relatively long compared to its period. Acoustic musical instruments typically have a large number of musical resonances. We show that for relatively small loop gains for collocated, linear control of an acoustic musical instrument, the single resonance model of Section 2.2.1 predicts the general behavior of musical resonances subject to control. This is a consequence of an acoustic musical instrument being a passive mechanical (positive real) system. In the computation of the departure angles of the poles associated with a musical resonance, it turns out that the departure angles are largely independent of the other musical resonances because their departure angle contributions cancel out approximately. This new result implies that we can use a simple model to predict the complex behavior of an acoustic musical instrument when subject to collocated, linear feedback control with relatively small gains.

Control Power For Changing A Resonance Frequency (Section 3.1.2)

A number of effects limit the efficacy of feedback controllers for acoustic musical instruments in practice. Consider controlling a single musical resonance with quality factor Q . We show that if the amount of control power is held constant, the maximum decrease in the decay time of a musical resonance is about $2Q$ times larger than the maximum increase in resonance frequency. In other words, a relatively large amount of control power is required for changing the resonance frequency of musical resonances using collocated linear control. As a consequence, we argue that, if a controller is to significantly adjust the fundamental frequency of an acoustic musical instrument, the controller should be able to directly actuate the wave speed of the medium.

Termination-Based Control and Matched Termination Control

We introduce and analyze a number of novel schemes for changing the fundamental frequency of a waveguide without actuating the wave speed directly. These schemes require a relatively large amount of control power (see previous point). *Termination-based control* requires only a single sensor and actuator. They can be placed near a termination of the waveguide. Controllers are designed by requesting specific waveguide termination reflectances (Section 3.3.2). In contrast, in matched termination control, the waveguide has a matched termination at one end, which allows waves to dissipate without reflection. A controller is designed which synthesizes a returning wave to induce musical resonances in the waveguide (Section 3.3.4). Unfortunately, it is difficult in practice to perfectly terminate an acoustic waveguide.

Control Bandwidth Limitations (Section 3.3.3)

Besides relatively large amounts of control power, relatively large loop gains are required for termination-based control. We argue that control bandwidth limitations prevent us from achieving such large loop gains at high enough frequencies in our laboratory. This implies that we cannot implement termination-based control of a plucked string in our laboratory without making significant improvements to our sensors, actuators, and digital controller.

Limitations On Non-Collocated Control Due To Sensor Nonlinearity (Sections 3.3.1 and 3.3.4)

In non-collocated control, there is not a sensor placed at the same location as every actuator. One of the inherent difficulties of non-collocated control is that the control system cannot typically rely strongly on any self-correction due to the time delays between sensors and actuators. We argue that, as a consequence, sensor nonlinearity (see Appendix F) can destabilize non-collocated control of musical resonances. Accordingly, sensors would have to have very low harmonic distortion in order to be employed successfully for non-collocated control of a plucked string. Since matched termination control is non-collocated, it is limited in the same sense by sensor nonlinearity. We hypothesize that Guérard was able to implement non-collocated traveling wave-based control of a recorder instrument because it has resonances with relatively short decay times [73].

Event-Based Control (Section 4.3)

In this context, self-sustaining oscillations describe oscillations in an acoustic musical instrument that sustain themselves. In event-based control, the controller repeatedly waits until an event has been detected and then actuates the acoustic musical instrument with a predetermined control signal in response. For example, when a pulse is injected into an acoustic waveguide-based instrument, it will eventually arrive at a sensor, which can cause another pulse to be triggered. We provide an analysis of how event-based control can be employed to specify the spectral envelope of the self-sustaining oscillations. In contrast with the aforementioned control techniques, event-based control offers a unique advantage for digital control configurations. It can shape the bandwidth of the musical instrument vibrations up to half of the sampling rate, even if there are anti-imaging filters.

RMS Level-Tracking Controllers (Section 4.4)

We introduce the formalization of RMS level-tracking controllers. These controllers monitor the RMS level of a signal in the feedback loop and use this information to

adjust control gains in real time. RMS level-tracking controllers are employed for implementing self-sustaining control of a waveguide using a variety of techniques. An important special case of these controllers is the dynamic range limiter. A dynamic range limiter can generate a control force signal by receiving state information that is processed by an audio effect.

1.5.2 Haptic Musical Instruments

Enabling The Musician To Make New Gestures (Section 6.3.2)

Feedback control can be employed to exert control over the gestures that a musician makes. For example, we demonstrate that feedback control in haptic musical instruments can enable the musician to make new gestures that would otherwise be difficult or impossible. To do so, we introduce two novel percussion instruments. The Haptic Drumstick is a haptically augmented drumstick that strikes a *virtual* surface, and the Haptic Drum is a haptically augmented but entirely *physical* drum pad. For both instruments, the drumstick is actuated with a pulse of energy each time that it strikes the drumming surface, enabling the musician to single-handedly play drum rolls. In contrast with traditional drumming interfaces, the Haptic Drum allows rolls to be played at speeds as fast as 70Hz. The feedback controller can adjust the strength of the pulses to alter the complexity of the drum rolls.

Dynamics Of The Haptic Drum (Section 6.2.4)

The haptic drum control algorithm aids the musician in playing drum rolls by promoting the formation of self-sustaining oscillations in the drumstick. These oscillations are known as limit cycles in the field of feedback control. We introduce a simple mechanical model describing the impedance of the musician's hand and the drumstick, and we employ a Poincarè map to study the dynamics of the haptic drum. We show that if each actuated pulse has the same intensity, then the limit cycle is stable.

Assisting Musicians In Making Gestures More Accurately (Chapter 7)

Feedback control, and haptic assistance in particular, can be employed to help musicians make gestures more accurately. O’Modhrain previously demonstrated that any haptic feedback was better than none at all for a simple pitch selection task [120]. We demonstrate that haptic assistance that is specific to the task can help further increase accuracy in comparison with more general kinds of haptic feedback. In particular, we believe it is important that the musician can regulate the degree of haptic assistance in real-time. The haptic assistance condition that we call FRC, for force-sensitive detents, provides detent-like feedback, where the strength of the detents is regulated by the downward pressure that the musician applies. It aids musicians in selecting pitches accurately when they press downward with at least some mildly large force. However, it does not prevent the musician from playing large intervals or glissandi when pressing downward only lightly.

Assistive Haptic Interfaces (Section 8.6)

Assistive haptic interfaces aid a user in completing tasks. Relying on prior research as well as our own, we provide guidelines for designing assistive haptic interfaces:

1. The most effective type of haptic assistance depends on the task (Section 7.8).
2. It can be helpful to allow the user to regulate the degree of haptic assistance in real time (Section 7.7.5).
3. If haptic assistance is deterministic and relatively simple, then it is more likely to be useful (Section 5.4.6).
4. Effective haptic assistance should never take an unexpected action so fast that it surprises the user, which could prevent him or her from responding appropriately (Section 5.4.6).
5. Many useful haptic assistance systems are not passive (Section 8.4).

Interfacing A Haptic Device With A Virtual, Digital-Waveguide Based Musical Instrument (Appendix A)

Other researchers have previously designed haptic musical instruments that mimic real acoustic instruments. The feedback controllers employed for this have been implemented using finite difference and (approximate) hybrid digital waveguide models. We present a novel method for constructing haptic musical instruments in which a haptic device is directly interfaced with a conventional digital waveguide model by way of a junction element, improving the quality of the musician's interaction with the virtual instrument. We prefer digital waveguide models because only a small number of multiply-adds are required, large portions of memory can be factored into simple digital delay lines, digital waveguide models are very accurate, and there is an extensive literature on calibrating them for realistic sound.

Chapter 2

Background: Feedback Controlled Acoustic Musical Instruments

2.1 Outline

2.1.1 Overview

Traditionally, the typical musician has owned many different kinds of instruments. The advantage has been that he or she could achieve a certain characteristic sound by playing the corresponding instrument. However, transporting the multitude of instruments to performances has been onerous, and musicians have required large amounts of space to store all of the instruments.

Rather than requiring different instruments in order to achieve different characteristic sounds, a musician might prefer to have a single instrument whose acoustics are programmable. To achieve different or new sound characteristics, the musician could simply load a different instrument program. Similarly, to construct a novel instrument, the musician could write a new “acoustical” program rather than going through the time-consuming process of mechanically constructing a new instrument from scratch. Furthermore, besides obviating the need to store large numbers of musical instruments, the interface to the instrument would remain similar despite changes in the acoustical behavior. This *interface constancy* would alleviate the musician

of the time-consuming process of learning to play via different interfaces merely to obtain different characteristic sound qualities.

Finally, acoustic musical instrument interfaces are tangible and allow the musician to develop a special kind of intimacy. Perhaps interface constancy combined with a wide array of possible acoustic-like instrument sounds present the strongest arguments for the utility of a feedback controlled acoustic musical instrument.

In the first half of this thesis, we study the problem of making the acoustics of an acoustic musical instrument programmable by way of feedback control. We could design a series of analog controllers, but digital control allows more freedom in controller implementation, so we design digital controllers. Nevertheless, in many instances, the digital controller approximates physical behaviors, so it is more convenient to mathematically derive analog controllers. The discretization from analog to digital controller form is equivalent to numerous well-studied filter design problems [122]. Hence, we do not specifically study the discretization problem in this thesis. Readers interested in the implementation details may consult the source code for the digital controllers (see Appendix G).

A feedback controlled musical instrument bears some resemblance to haptic musical instruments in that a feedback controller affects the interaction between the musician and the instrument [41]. However, feedback controlled acoustic musical instruments remain closer to the real world than the virtual world because the majority of the instrument is real rather than virtual. A feedback controlled acoustic musical instrument might nonetheless be considered a haptic musical instrument whose haptic interface is the whole acoustical medium.

2.1.2 Musical Acoustics Review

The field of musical acoustics for traditional acoustic musical instruments is quite mature, so we cannot hope to cover it here in any detail [62]. However, to ensure that this document remains accessible to readers in other fields, particularly electrical, mechanical, and aeronautical engineering, we briefly mention a few points here.

Many acoustic musical instruments can be described as nearly linear vibrating

structures, although some important exceptions are outlined in Chapter 4. Consequently, the usual modal analysis techniques are useful for studying the modal shapes, and eigenvalues [62]. In contrast with many other vibrating structures, especially the lower frequency modes of musical instruments tend to have long decay times on the order of seconds (see Table 3.1). In addition, a relatively large number of modes (e.g. > 20) may participate significantly in typical vibrations. These facts are commensurate with the observation that while most structural engineers seek to prevent structural vibrations, instrument builders seek to create sustained structural vibrations.

A large number of acoustic musical instruments are harmonic. Then the significant resonant frequencies of vibration are a subset of the harmonic series $f_0, 2f_0, 3f_0, \dots$. The fundamental frequency f_0 is the lowest significant resonant frequency. However, the *pitch* is defined to be the listener's perception of the fundamental frequency. In many instruments, the modal frequencies may become slightly misaligned or inharmonic. In this case, a listener will often perceive the pitch to be different than f_0 . In more complicated situations, the concept of pitch may even be ambiguous: different listeners will perceive significantly different pitches [59].

2.1.3 Problem Description

Figure 2.1 depicts how a musician interacts with a feedback controlled acoustic musical instrument [28]. The upper part of the block diagram shows the same type of interaction that arises with a traditional acoustic musical instrument. The musician hears the acoustic output from the instrument and feels the vibrations of the instrument. In response, the musician applies an excitation force to the instrument. For example, the musician might blow air across a reed or pluck a vibrating string. Over time, the musician also alters parameters of the musical instrument, such as the fundamental frequency.

The lower part of the block diagram in Figure 2.1 depicts how feedback control changes the acoustical properties of the instrument. v represents the controller's observation of the instrument state, such as the velocity. The control output signal

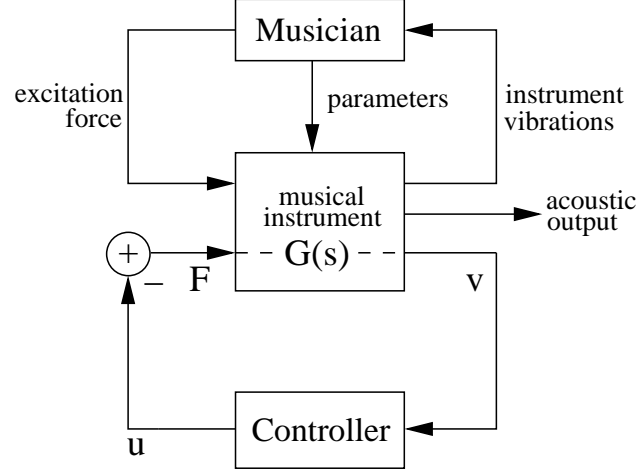


Figure 2.1: Musician interacting with a feedback controlled acoustic musical instrument

from the controller u is negated and applied as the force F to the instrument by an actuator. $G(s)$ is the mechanical mobility of the instrument as seen from the sensor and actuator.

Figure 2.2 shows a simplified version of Figure 2.1, which focuses on the role of the controller. r represents the excitation force exerted by the musician. The form of the diagram is perhaps slightly misleading in that different modes may be controllable from the point of view of the musician and the actuator. For instance, if the actuator is placed on the node of a certain mode, feedback control will not be able to affect this mode. These controllability differences between the actuator excitation and the musician's excitation, however, turn out to be of minor importance in the upcoming development of the controllers, chiefly because many of the analyses assume that $r = 0$.

The musical instrument's dynamics usually change with time according to parameters dictated by the musician, such as fundamental frequency. Ideally the closed loop performance should not be too sensitive to instrument parameters such as fundamental frequency, as the musician may vary these over time. This requirement may seem overly restrictive, but we will still show that many controllers can be made insensitive to the resonance frequencies of the instrument modes. For other controllers that are




Figure 2.2: Simple block diagram for feedback control of an acoustic musical instrument

sensitive to the instrument's modal frequencies, we suggest augmenting the instrument with sensors such that these parameters can be sensed accurately in real-time without being disturbed by the acoustic feedback loop.

2.1.4 Challenges

- Perhaps the foremost challenge lies in designing useful controllers that are minimally dependent on $G(s)$, which changes as a function of the fundamental frequency. We choose this approach because estimating the fundamental frequency of an acoustic musical instrument can be nontrivial, especially when the estimate needs to be calculated quickly [132].
- Control engineers often employ one sensor per mode to be controlled. Since most musical instruments have a large number of significant modes, applying one sensor per mode becomes unrealistic. In addition, feedback control applications typically involve adjusting a small number of significant resonances. However, the large number of resonances in typical acoustic musical instruments can complicate controller design, especially when applying methods such as pole placement [65].
- The modes for acoustic musical instruments typically have long decay times (see Table 3.1), and in many cases we wish to preserve the long decay times under control to create musical dynamics. Hence perturbations and nonidealities in

the feedback controller, plant, sensors, and/or actuators are more likely to result in unstable closed loop behavior.

- There are no commercially available sensors or actuators that are optimal for feedback control implementations. We need to design and construct our own sensors and actuators. These designs need to minimize non-ideal effects such as mismatching, noise, and nonlinearity.

2.1.5 Prior Examples

Various forms of actively controlled musical instruments have already been designed. For instance, musicians have investigated the problem of indefinitely sustaining string vibration, which amounts to inducing a stable limit cycle in a vibrating string so that a single note lasts as long as desired. To this end, musicians have used acoustic feedback from power amplifiers to actuate their electric guitar strings in an ad-hoc manner; however, due to acoustic delays and the nonlinear nature of the amplifiers, this approach has been difficult to control precisely. The commercially available EBow [82] and Sustainiac [84] have mitigated this problem using electromagnetic actuators. Similarly, Weinreich and Caussé have electromagnetically induced the Helmholtz “stick-slip” bowing motion in a vibrating string [168].

Active control has been applied to the body structures of acoustic guitars. For example, in order to suppress tendencies toward acoustic feedback in amplified situations, Griffin actively damped the first plate mode in an acoustic guitar [71]. Hanagud boosted the second plate mode in an inexpensive acoustic guitar to make it sound more like an expensive acoustic guitar [78]. Chollet, Aeberli, and Charles Besnainou actively tuned the Helmholtz body resonant frequency of a guitar [48].

Besnainou also applied active control to a violin, a snare drum, a pipe organ, and a marimba bar [28]. Boutin has changed the damping time, fundamental frequency, and gain of a xylophone bar using Proportional-Integral-Derivative (PID) control [33][34]. Boutin has also applied PID control to control the physical behavior of a violin bridge [35]. Other researchers have actively controlled drums as well [92][163].




Figure 2.3: Damped harmonic oscillator with mass m , damping parameter R , spring constant k , and external force $F(t)$

2.2 Basic Models

2.2.1 Single Resonance Model Of Instrument

In the following chapters, we attempt to analyze each controller in the framework of the simplest applicable model. The simplest possible model of a musical instrument is that of a single resonance. For instance, a mass m , damper with damping constant R , and a spring with stiffness k sharing the same velocity form a damped harmonic oscillator as shown in Figure 2.3. If an external force F is applied to the oscillator, then a differential equation describing the dynamics of the forced, damped harmonic oscillator may be written:

$$m\ddot{x} + R\dot{x} + kx = F = -u \quad (2.1)$$

The decay time τ for a musical oscillator is generally relatively long.

$$\tau = 2m/R \quad (2.2)$$

See Table 3.1 on page 82 for some sample values. If the decay time is long, then the oscillator is considered lightly damped, and the fundamental frequency

$$f_0 \approx \frac{1}{2\pi} \sqrt{\frac{k}{m}}. \quad (2.3)$$




Figure 2.4: Digital waveguide model of a vibrating string with one sensor and one actuator

2.2.2 Waveguide Model

In other cases, it is important to consider all of the instrument modes. Then we employ a waveguide model, which models vibrations of a one-dimensional waveguide with inverting reflections at the terminations [149]. For most controllers, we have one sensor and one actuator, so this is the case we consider for the digital waveguide model in Figure 2.4. The model can easily be extended to have multiple sensors and actuators using the superposition principle.

We use the delay element z^{-1} to model a propagation delay of $1/f_s$ seconds, where f_s is a sampling rate. Multiple delay elements can be combined to model longer delays [149]. The top three delay lines $z^{-S/2}$, $z^{-A/2}$, and $z^{-(N-A-S)/2}$ model waves traveling to the right, and the bottom three delay lines $z^{-(N-A-S)/2}$, $z^{-A/2}$, and $z^{-S/2}$ model waves traveling to the left. f_s/N is approximately equal to the fundamental frequency f_0 . Figure 2.5 shows the proportionality relationship between S , A , and N and sensor distance, sensor-to-actuator distance, and the total string length, respectively. $H_{lp}(z)$ is a linear-phase low-pass filter that causes higher resonances to decay more quickly. If $H_{lp}(z)$ were not linear phase, then the digital waveguide model would model a vibrating string with dispersion [62].

Now we make the model specific to displacement waves in a vibrating string, although the model could be easily adapted to many other one-dimensional waveguide instruments [149]. The displacement $x[n]$ measured by the sensor (see Figure 2.4) is the sum of the two traveling displacement waves at the sensor position. The traveling waves change sign at the string ends due to the boundary condition that the string displacement is zero at the ends. $\Delta x_{act}[n]$ is the change in displacement due to the




Figure 2.5: Relative dimensions of string model

actuator. It may be computed by digitally integrating the change in velocity $\Delta v_{act}[n]$. Finally, since $F[n]$ is the force signal to be actuated,

$$\Delta v_{act}[n] = \frac{F[n]}{2R_0}, \quad (2.4)$$

where R_0 is the string's wave impedance [149].

2.3 Passive Control

2.3.1 Definition

On average, passive controllers do not add any energy to the instrument. For this reason, they are guaranteed to result in stable closed-loop behavior, irrespective of the instrument [138]. This property implies that the exact modal frequencies and decay times of the uncontrolled instrument do not affect the closed loop stability as long as the instrument is dissipative—hence passive controllers will be stable regardless of the instrument's fundamental frequency. The musician can change notes without worrying that the system may become unstable.

For similar reasons, passive control plays a key role in haptics, where a human interacts directly with actuators and sensors. The human's impedance may vary over time, so haptic controllers are often passive, ensuring theoretically that the system remains stable [56]. Practical considerations always prevent controllers from being completely passive, so the analysis of closed loop behavior under extreme conditions becomes more complicated (see Section 2.5.3 and [56]). Nevertheless, the positive real formulation accurately predicts behavior in many situations. In the next few sections, we introduce the positive real formulation of passivity for linear systems, and we provide a stability proof for feedback control of a dissipative acoustic musical

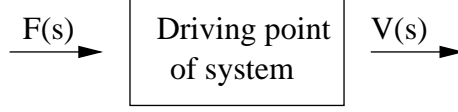


Figure 2.6: Driving point of a system

instrument under positive real control.

2.3.2 Positive Real Functions

Before defining positive real functions, we first introduce the Laplace transform. Let $\mathcal{L}\{g\} = G$ be short hand for the following definition of the unilateral Laplace transform [65]:

$$\mathcal{L}\{g(t)\} = G(s) = \int_{0^-}^{\infty} g(t)e^{-st}dt. \quad (2.5)$$

$g(t)$ is the impulse response of the acoustic musical instrument, and $G(s)$ is the Laplace domain representation of the impulse response [65]. Often the complex variable $s = j\omega + \sigma$ is decomposed into the real part σ and the imaginary part $j\omega$.

Figure 2.6 depicts the concept of the driving point of a system. If the force $F(s)$ is exerted on a system at a single point and the velocity $V(s)$ is measured at the same point, then the *driving point impedance* is defined to be $\frac{F(s)}{V(s)}$. The *driving point mobility* is defined to be $\frac{V(s)}{F(s)}$.

A rational function $G(s)$ is *positive real* if [161]

1. $G(s)$ is real when s is real.
2. $\operatorname{Re}\{G(s)\} \geq 0$ for all s such that $\operatorname{Re}\{s\} \geq 0$.

A rational function $L(s)$ is *strictly positive real* if

1. $L(s + \epsilon)$ is positive real for all real $\epsilon > 0$.

Now we note some of the properties of positive real and strictly positive real functions [161][149]:¹

¹Note: The bilinear transform preserves s-domain and z-domain sense positive realness [149].

1. $|\angle G(j\omega)| \leq \frac{\pi}{2}$ for all ω .
2. $|\angle L(j\omega)| < \frac{\pi}{2}$ for all ω .
3. $1/G(s)$ is positive real.
4. $1/L(s)$ is strictly positive real.
5. If $G(s)$ represents either the driving point impedance or driving point mobility of a system, then the system is *passive* as seen from the driving point.
6. If $L(s)$ represents either the driving point impedance or driving point mobility of a system, then the system is *dissipative* as seen from the driving point.
7. $G(s)$ and $L(s)$ are stable.
8. $G(s)$ and $L(s)$ are minimum phase.
9. The relative degrees of $G(s)$ and $L(s)$ must be less than 2.
10. If a point sensor and a point actuator operate on a structure and are collocated, then both $V(s)/F(s)$ and $F(s)/V(s)$ are strictly positive real because there must be some damping at all frequencies.
11. If a matched sensor and actuator operate on a structure and are collocated, then both $V(s)/F(s)$ and $F(s)/V(s)$ are strictly positive real because there must be some damping at all frequencies (see Section 2.5).
12. No matter what causal time-domain function $f(t)$ is used to excite the driving point, the velocity response $v(t)$ will be such that $\int_0^\infty f(t)v(t)dt \geq 0$.

2.4 Control System Stability

The controller has the ability to continuously add energy to the acoustic musical instrument since the controller is active and has a power supply. This characteristic implies that the control system is capable of becoming unstable, which is highly

undesirable. In the following, we provide and prove a sufficient condition for the stability of the control system assuming that the controller is linear.

2.4.1 Revised Bode Stability Criterion

Many engineers are familiar with the Bode stability criterion, which is commonly used for analyzing the stability of simple systems [65]. However, the criterion is a sufficient, but not necessary condition for system instability. For this reason, we present the less well-known but more general “revised Bode stability criterion” [76]:

A closed-loop system is stable if the open-loop system is stable and the frequency response of the open-loop transfer function has an amplitude ratio of less than unity at all frequencies corresponding to an angle of $-\pi - n(2\pi)$, where $n = 0, 1, 2, \dots, \infty$.

The revised Bode stability criterion can be derived from the Nyquist criterion by observing that the above statement means that no encirclements of the $(-1,0)$ point can exist on the Nyquist diagram.

2.4.2 Stability Proof

In this section, we employ the revised Bode stability criterion to argue that the control system must be stable if the controller $K(s)$ is positive real and the acoustic musical instrument model $L(s)$ is strictly positive real. Due to property 7 from Section 2.3.2, $K(s)$ and $L(s)$ are stable, and so the open-loop system $K(s)L(s)$ is stable. The criterion requires only a further analysis along the frequency axis, so we consider the transfer functions evaluated at $s = j\omega$. Figure 2.7 serves as a reminder of the configuration.

Properties 1 and 2 from Section 2.3.2 imply that

$$|\angle(L(j\omega))| + |\angle(K(j\omega))| < \pi \quad \forall \omega, \quad (2.6)$$

so the absolute value of the angle of the products must also be upper bounded by π

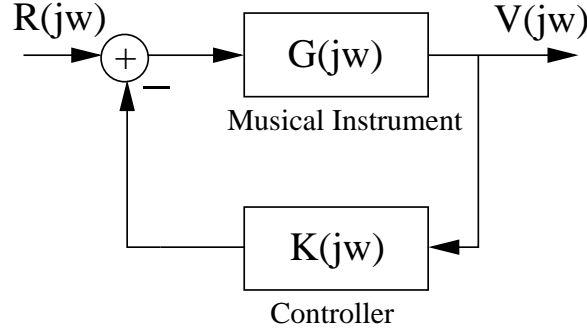


Figure 2.7: Simple block diagram of control system with $s = j\omega$

[147]:

$$|\angle(L(j\omega)K(j\omega))| < \pi \quad \forall \omega. \quad (2.7)$$

Consequently, the revised Bode stability criterion implies that the control system is stable.

2.4.3 Implications

In summary, if $L(s)$ accurately models the acoustic musical instrument and is strictly positive real, then any positive real controller $K(s)$ applied to it results in a stable closed-loop system. Furthermore, the proof still holds if we change the controller to $K_0K(s)$ for some real constant $K \geq 0$, implying that the system remains stable even for infinitely large loop gains. This amazing property is known as unconditional stability [154][10], and it implies that the control system will be endowed with a large amount of control authority over the acoustic musical instrument.

In summary, the control system is unconditionally stable, regardless of the particular parameters of the acoustic musical instrument, if

1. the sensors and actuators are collocated and matched and the acoustic musical instrument is dissipative (implying that $G(s)$ is strictly positive real) and
2. the controller $K(s)$ is positive real.

The most important consequence here is that when designing a feedback controlled musical instrument, the sensors and actuators should be collocated, and they should

be matched as well as possible.

2.5 Practical Collocation Considerations

The results from the previous section are quite impressive; however, they are only theoretical. In practice there are several effects that can prevent $G(s)^2$ from being perfectly positive real, thus limiting the maximum stable loop gain. Besides noncollocation and mismatching effects, sensor and actuator nonlinearity, inductive coupling between sensors and actuators (see Appendix F), and controller delay can limit the maximum stable loop gain.

2.5.1 Matched Sensor and Actuator

Neither a real sensor nor a real actuator operates precisely at a point. Instead, a real sensor and a real actuator operate over finitely sized spatial sensitivities of the medium being controlled. If the sensor and actuator geometry are the same, in other words, if they both operate over the same spatial sensitivity, then $G(s)$ is still positive real, and we can make use of the convenient properties described in Section 2.4. An example of a perfectly matched sensor and actuator would be a pair of identical ideal piezoelectric bending patches on either side of an infinitely thin membrane. With the membrane held horizontally, the patch on top of the membrane would be placed directly over the patch on the underside of the membrane [94].

Marty Johnson and Stephen Elliott proved that $G(s)$ is positive real if a *sensor and an actuator are matched*. They used a modal decomposition, which we extend here to our variable and resonance model conventions [94]. Consider a surface that vibrates transversely at position (x, y) according to the complex Laplace variable s with velocity $\tilde{V}(x, y, s)$. The velocity can be decomposed into the sum of the (real) modal shapes $\phi_n(x, y)$ multiplied by the (complex) amplitudes $a_n(s)$, where n denotes

²We choose to lump the sensor and actuator nonidealities into $G(s)$ rather than the controller transfer function $K(s)$.

the modal number:

$$\tilde{V}(x, y, s) = \sum_{n=1}^{\infty} a_n(s) \phi_n(x, y). \quad (2.8)$$

We assume that the sensor and the actuator both operate with the spatial sensitivity $\eta(x, y)$, then $a_n(s)$ can be written as follows:

$$a_n(s) = F(s) A_n(s) \int \int \phi_n(x, y) \eta(x, y) dx dy, \quad (2.9)$$

where $F(s)$ is the control force input and $A_n(s)$ is a second-order resonance term. For instance, if we use the multi-resonance model given in Section 3.1.3 and we find equivalent masses $m_n > 0$, damping parameters $R_n \geq 0$, stiffnesses $k_n > 0$, and scaling constants $\gamma_n > 0$, then

$$A_n(s) = \frac{\gamma_n s}{m_n s^2 + R_n s + k_n}, \quad (2.10)$$

analogously to (3.25). Note that $A_n(s)$ is positive real. Now let the sensor measure velocity $V(s)$.

$$V(s) = \int \int \tilde{V}(x, y, s) \eta(x, y) dx dy \quad (2.11)$$

By applying (2.8) and (2.9), we can write

$$V(s) = F(s) \sum_{n=1}^{\infty} A_n(s) \left[\int \int \phi_n(x, y) \eta(x, y) dx dy \right]^2. \quad (2.12)$$

As a consequence, we can write the mechanical “driving surface” mobility $G(s)$ of the musical instrument as seen by the controller as the following:

$$G(s) = \frac{V(s)}{F(s)} = \sum_{n=1}^{\infty} A_n(s) C_n, \quad (2.13)$$

where each $C_n \geq 0$ is given by the brackets-squared term in (2.12). Since $A_n(s)$ is positive real and $C_n \geq 0$, (2.13) implies that $G(s)$ must also be positive real. On the other hand, if the sensor and the actuator do not share the same spatial sensitivity $\eta(x, y)$, then some of the C_n will become negative for some n , possibly causing $G(s)$

to no longer be positive real. For the cases involving a vibrating string that we have investigated numerically, any mismatch in the spatial sensitivities prevented $G(s)$ from being positive real [21]. An extreme case is when the sensor and the actuator spatial sensitivities do not overlap at all. Then there is a pure delay in between them, preventing $G(s)$ from being positive real as explained in Section 2.5.3.

Nevertheless, collocating a sensor and an actuator in practice in other contexts can be challenging. First of all, the geometric constraints can be prohibitive. In addition, placing a sensor and an actuator in close proximity to one another may result in direct inductive coupling (see Section F.4). For more details on the non-ideal characteristics of our guitar string sensor and actuator in the laboratory, please see Section 2.1.9 of Technical Report STAN-M-120 [21].

2.5.2 Sensoriactuators

A sensoriactuator is a single transducer device that is concurrently employed as both a sensor and an actuator. If the sensoriactuator operates ideally, then the sensor and actuator are perfectly collocated and matched. Sensoriactuators have been studied since 1961 at the latest for feedback control of loudspeakers [103][55]. Their application to piezoelectric transducers has been studied as well [6][75][4].

Unfortunately, perfect sensoriactuators tend to be difficult to realize in practice. One of the chief problems is that the same pair of wires is used for sensing and actuation, so it is difficult to electrically isolate the sensing and actuation signals from each other. Typically a bridge circuit is used to provide a subtraction signal from a model of the transducer; however, effects such as temperature drift, hysteresis, and other nonlinearities complicate the implementation of the model. Efforts at adaptive compensation of sensoriactuators have attempted to overcome these limitations although only with limited success [49].

Ierymenko has investigated using time multiplexing to implement an approximate electromagnetic sensoriactuator. At each instant in time, the pair of wires is either used to communicate a sensing signal or an actuation signal [90]. This configuration appears to solve many of the problems described above, but is nevertheless subject

to a control system delay, whose nonideal effects are described below. Some researchers have applied a related technique to controlling switched reluctance motors [88][133][12][9].

2.5.3 Digital Control System Delay

When a digital controller is employed to close the feedback loop, there will always be some controller delay of D seconds. The delay is due to the analog-to-digital converter (ADC), processing time, the digital-to-analog converter (DAC), and output signal reconstruction in the digital controller (see Appendix G).

With delay D , the controller transfer function becomes $K(s)e^{-sD}$. The delay contributes to significant phase lags at higher frequencies, preventing the controller from implementing a truly positive real impedance in practice. For example, at infinite frequency, the phase delay is ∞ even if the desired $K(s)$ is positive real.



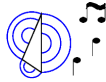
$$\lim_{\omega \rightarrow \infty} \angle(K(j\omega)e^{-j\omega D}) = -\infty \quad (2.14)$$

This phase lag can limit the maximum achievable loop gain, which limits the degree by which the control system can affect the musical instrument. Consequently, if especially large loop gains are required, a digital delay can prevent them from being achievable even if $1/D$ is much larger than the system bandwidth (see Section 3.3.3). This property implies that depending on the particular control application, control system designers should consider implementing analog controllers instead.

2.6 Organization

In chapters 3 and 4, we develop both linear controllers and self-sustaining controllers with application to acoustic musical instruments. Since the passivity-based controllers apply to all dissipative acoustic musical instruments, we discuss them in the most general context. However, we test our controllers in the laboratory on a vibrating string, so it is sometimes more convenient to consider controllers specific to one-dimensional waveguides or even vibrating strings themselves. To make this document

Table 2.1: Classification of acoustic musical instrument types

	Stringed instruments
	One-dimensional waveguide instruments
	Dissipative acoustic musical instruments

as readable as possible, we place an icon at the beginning of each section of controllers to identify to which instruments the controller derivations and discussions apply.

Table 2.1 outlines the meaning of the icons. The class of stringed instruments, the most specific class, lies at the top of the table. Examples of the more general class of one-dimensional waveguide instruments include not only stringed instruments, but also clarinets, flutes, oboes, trumpets, etc. Finally, drums, cymbals, and other instruments with more complicated dynamics belong to only the most general class of dissipative acoustic musical instruments, which is placed at the bottom of Table 2.1.

Figure 2.8 provides an overview of the controller classes covered in chapters 3 and 4. Positive real controllers, the first class, are stable in theory for any dissipative acoustic instrument $G(s)$; however, they may not be applied easily in practice to change the fundamental frequency of the instrument. In contrast, non-collocated traveling wave-based controllers can control the fundamental frequency of wind and bowed-string instruments, but they are sensitive to sensor nonlinearities. Furthermore, arrays of sensors are required for estimating the traveling wave states. Termination-based controllers are also capable of changing the fundamental frequency, but relatively large loop gains are required, preventing these controllers from being applied to plucked string instruments in practice. This is unfortunate because the termination is the only point along the string where all of the modes are controllable and observable.

As discussed in Chapter 4, passive nonlinear controllers are guaranteed to be stable when applied to dissipative musical instruments (see Figure 2.8). Self-sustaining control may be applied to sustain the vibrations of a musical instrument for an arbitrarily long time. The self-sustaining characteristic allows the musician to concentrate on shaping the timbre of the sound rather than concentrating on preventing vibrations from decaying.





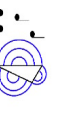

Class	Analysis Class	Linear?	Self-Sustaining?	Passive?	Mechanical Analog?	Useful For	Pros	Cons	Implemented?	Simulated?
Positive Real		Y	N	Y	Exists	Controlling musical instruments when details about the dynamics of the instrument (e.g. fundamental frequency) are unknown	In theory, guaranteed to be stable as long as the instrument $G(s)$ is dissipative.	It is more difficult in practice to obtain specific behaviors, such as increasing the fundamental frequency of an instrument.	Y	Y
Non-Collocated Traveling Wave-Based		Y	N	N	Does not exist	Controlling the damping and fundamental frequency of wind and bowed-string instruments	In theory, controller design is analogous to digital waveguide virtual instrument design.	Sensitivity to sensor nonlinear prevents applications in instruments with naturally long decay times such as plucked-stringed instruments. Arrays of sensors are required.	N	Y
Termination-Based		Y	N	N	Does not exist in general	Controlling the damping and fundamental frequency of 1-D waveguide instruments	Only one sensor is required.	Sensitive to sensor nonlinearity (as above). Relatively large loop gains are required in practical contexts.	N	Y
Matched Termination Control		Y	N	N	Does not exist in general	Controlling the damping and fundamental frequency of 1-D waveguide instruments	Only one sensor is required.	Sensitive to sensor nonlinearity (as above). It is difficult in practice to perfectly terminate an acoustic waveguide.	N	N
Passive Nonlinear		N	N	Y	Exists	Implementing nonlinear damping and stiffness modulation effects	Stability is guaranteed by passivity. Implements nonlinear modal coupling.	More musically interesting self-sustaining nonlinearities are not passive, so they cannot be implemented.	Y	Y
Self-Sustaining Control		N	M	N	Does not exist in general	Sustaining the vibrations of a musical instrument for an arbitrarily long time	Energy arguments show that the controlled system cannot become unstable.	Analysis of nonlinear systems can be difficult.	Y	Y


Figure 2.8: Overview of controller classes covered in this document relating to the feedback control of acoustic musical instruments

Chapter 3

Changing the Linear Dynamical Properties of Acoustic Musical Instruments

Linear control is most useful for changing the linear properties of acoustic musical instruments. Besides altering the damping of modes, it may for example be applied to change the resonant frequencies of modes. In this chapter, we first describe positive real linear controllers. Then we move on to discuss a few lower order controllers that are not positive real as well as higher order controllers and their limitations.

3.1 Passive Control



Passive control is important because in theory, a passive controller cannot drive a dissipative acoustic musical instrument unstable, no matter how high the loop gain is. This statement implies that passive controllers have considerable control authority over acoustic musical instruments because relatively large loop gains can be applied even in practice. Figure 3.1 shows an example configuration for controlling a vibrating string with a single sensor, which is depicted as a microphone, and a single actuator, which is depicted as a loudspeaker. In order to implement passive control,




Figure 3.1: Collocated control of a vibrating string with a single sensor and actuator

the sensors and actuators must be collocated (see Section 2.4). Consequently, the sensor and the actuator in Figure 3.1 are shown to operate at the same point along the string.

Passive linear controllers satisfy the positive real property, which also implies that mechanical analogs exist for them (see Section 2.3.2). As a consequence, each passive linear controller has a mechanical interpretation. The low-order positive real controllers are particularly easy to implement in practice, so we describe them first, and we explain what effect they have on the closed-loop instrument behavior. Before we jump into all the details, we provide a summary in Figure 3.2 of the positive real controllers that we will investigate in Section 3.1. Each controller name corresponds to the name of a section in Chapter 3.

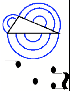

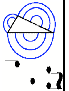

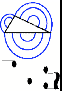


Name	Analysis Class	Linear?	Self-Sustaining?	Passive?	Mechanical Analog	Useful For	Pros	Cons	Implemented?	Simulated?
Proportional-Integral-Derivative (PID) Control		Y	N	Y	Damped harmonic resonator consisting of a mass, spring, and damper	Changing the overall damping; altering the damping and resonance frequency of the lowest resonance	Simplicity. Stable in practical situations for a large range of musical instrument parameters.	There are not enough control parameters to specify the damping times and resonance frequencies of the all of the musical instrument's modes.	Y	Y
Bandpass filter		Y	N	Y	Mass, spring, and damper physically in series	Shortening the decay time of a specific resonance	Primarily affects resonances with resonance frequencies similar to that of the controller itself.	Behavior depends on the resonances of the musical instrument.	Y	Y
Notch filter		Y	N	Y	Yes – see text	Shortening the decay times of all of the resonances except for a specific resonance	Simplicity.	Behavior depends on the resonances of the musical instrument.	Y	Y
Alternating filter		Y	N	Y	Exists	Shifting the fundamental frequency of a 1-D waveguide	The resonance frequencies of the modes are shifted by similar amounts.	Behavior depends on the resonances of the musical instrument.	N	Y
Teleoperator-based		Y	N	Y	Mass and spring physically in parallel	Imposing the dynamics of a model onto a physical musical instrument	The dynamics of the instrument are specified directly by the model.	Large loop gains are required. The mass, stiffness, and resonance frequencies of the physical instrument can cause detuning of the model parameters.	Y	Y
Gyrator		Y	N	Y	Does not exist	Detuning the resonant modes of a 1-D waveguide	Modes across the entire frequency range are significantly affected.	The decay times are also perturbed. Some become longer and others shorter, so for large loop gains, one mode dominates, decaying slowly.	N	Y

Figure 3.2: Summary of positive real controllers covered in Section 3.1

3.1.1 Proportional-Integral-Derivative (PID) Control



Proportional-Integral-Derivative (PID) control is employed in many applications because of its simplicity. The feedback force is a sum of the scaled versions of a measured variable (P), its integral (I), and its derivative (D) [65]. If we measure the velocity and the actuator is collocated with the sensor, then the controller does not change the model order. In fact, it allows the mass, damping parameter, and spring constant to be adjusted directly by the control variables. However, since we measure the displacement in the laboratory, we write the feedback equation (3.1) in terms of the displacement x instead. There is a term proportional to the displacement (P_P), a term proportional to the derivative of the displacement (P_D), and a term proportional to the acceleration of the displacement (P_{DD}):

$$u \triangleq P_{DD}\ddot{x} + P_D\dot{x} + P_Px. \quad (3.1)$$

The feedback control law can also be expressed in the Laplace domain assuming zero initial conditions, where $\dot{x}(t) \longleftrightarrow V(s)$:

$$K_{PID}(s) = \frac{U(s)}{V(s)} = P_{DD}s + P_D + P_P/s. \quad (3.2)$$

We analyze the dynamic behavior assuming that the acoustic musical instrument can be modeled as having a single mechanical resonance. We show in Section 3.1.3 that these results generalize to acoustic musical instruments with multiple modes if they all have long decay times. Substituting (3.1) into the dynamics of a single resonance plant (2.1), we obtain a differential equation describing the closed-loop system [28].

$$(m + P_{DD})\ddot{x} + (R + P_D)\dot{x} + (k + P_P)x = 0 \quad (3.3)$$

When the new equivalent mass, damping parameter, and spring constant are substituted into (2.2) and (2.3), we obtain the decay time constant with control $\hat{\tau}$

and the resonance frequency with control \hat{f}_0 :

$$\hat{\tau} = \frac{2(m + P_{DD})}{R + P_D} \quad (3.4)$$

$$\hat{f}_0 \approx \frac{1}{2\pi} \sqrt{\frac{k + P_P}{m + P_{DD}}} \quad (3.5)$$

Differentiating the displacement twice to obtain an estimate of the acceleration is often too noisy in practice [65], but we notice we can adjust $\hat{\tau}$ and \hat{f}_0 independently even if $P_{DD} = 0$. Then P_D adjusts the decay time, and P_P is useful for changing the resonance frequency.

In our laboratory, we apply the controller to a vibrating string, as depicted in Figure 3.3 (top), where $K(s)$ represents the controller. The microphone symbol represents the sensor, and the loudspeaker symbol represents the actuator. Figure 3.3 (bottom) illustrates the mechanical analog of the controller. In theory, the controller is equivalent to attaching a point mass P_{DD} , damper P_D , and spring P_P to the string at the point where the sensor and actuator are collocated. The existence of a mechanical analog is proof that PID control is positive real and therefore passive [161].

Given the positive real interpretation from Section 2.3.2, we know that velocity and force are fundamental variables from an energy perspective that describe structural vibrations. The control term adjusted by P_D relates these two quantities without memory, and is in this sense the most fundamental controller. Furthermore, the term P_D can only be responsible for damping. In fact it damps all of the controllable/observable modes simultaneously [154].

In contrast, the P_P term has a low-pass characteristic with respect to the P_D term. Therefore, P_P affects the low frequency modes more than the high frequency modes. In practice, the P_P term brings about a noticeable resonance frequency shift of primarily the lowest mode. Thus, P_P changes the pitch of the instrument, but it also detunes the harmonic series, which can be undesirable in some contexts. Considering the mechanical analog of P_P , the detuning is not a surprise. The normal practice of changing a vibrating string's fundamental frequency is to change its length, tension, or linear mass density, but not to attach a spring to it as in Figure 3.3 (bottom).




Figure 3.3: Collocated control of a vibrating string with a collocated sensor (depicted as microphone) and and actuator (depicted as loudspeaker) (top); Mechanical analog of PID control (bottom)

Higher order controllers are required to change the pitch of the instrument without detuning the harmonic series.

We have further simulated PID control for the digital waveguide model and made it available on the Internet in a tutorial format [25]. Although a single sample of delay is present in the controller model, this delay is similar in magnitude to the delay inherent in our hardware controller (see Section G).

3.1.2 Changing The Resonance Frequency Is Difficult



Consider a resonance with a long decay time relative to its period. We will call this a *musical* resonance. In this section, we show that changing the resonance frequency of a musical resonance requires significantly more control energy than changing the decay time. Let us consider controlling the resonance modeled by the mass, spring, and damper from Section 2.2.1. For a second-order system such as this, the quality

factor

$$Q \triangleq \frac{k}{R(2\pi f_0)} = \frac{m(2\pi f_0)}{R}, \quad (3.6)$$

where f_0 is the uncontrolled frequency of vibration. For musical resonances, the decay time is approximately Q periods [147]. **We will show that if the amount of control power is fixed, the maximum decrease in the decay time is about $2Q$ times larger than the maximum increase in resonance frequency!**

We present an example using displacement (P_P) and velocity (P_D) feedback to independently control the resonance frequency and the decay time. Without loss of generality, we assume an impulse $\delta(t)$ is added to the force signal to find the impulse response of the controlled system. In other words, the force signal $F(t)$ is the following:

$$F(t) = -P_D \dot{x}(t) - P_P x(t) - \delta(t), \quad (3.7)$$

where the first two terms are the feedback terms and the third term is due to an external excitation causing the system to vibrate. Consequently, the system begins to move with displacement $x(t) = -\frac{1}{m2\pi\hat{f}_0} e^{-t/\hat{\tau}} \sin(2\pi\hat{f}_0 t)$ for $t > 0$. The velocity is as follows:

$$\dot{x}(t) = \frac{1}{m\hat{\tau}2\pi\hat{f}_0} e^{-t/\hat{\tau}} \sin(2\pi\hat{f}_0 t) - \frac{1}{m} e^{-t/\hat{\tau}} \cos(2\pi\hat{f}_0 t); \quad (3.8)$$

however, the expression for the velocity can be simplified because the decay time is long relative to the period of vibration. In other words, $\hat{\tau}(2\pi\hat{f}_0) \gg 1$, and so for $t > 0$

$$\dot{x}(t) \approx -\frac{1}{m} e^{-t/\hat{\tau}} \cos(2\pi\hat{f}_0 t). \quad (3.9)$$

The feedback control signal $u(t)$, which does not include the impulse excitation $\delta(t)$, is

$$u(t) = P_D \dot{x}(t) + P_P x(t) \approx -\frac{P_D}{m} e^{-t/\hat{\tau}} \cos(2\pi\hat{f}_0 t) - \frac{P_P}{m2\pi\hat{f}_0} e^{-t/\hat{\tau}} \sin(2\pi\hat{f}_0 t). \quad (3.10)$$

Calculating The Control Energy

We calculate the unitless energy E of the control signal $u(t)$ where $A \triangleq -\frac{P_P}{m2\pi\hat{f}_0}$ and $B \triangleq -\frac{P_D}{m}$:

$$E = \int_0^\infty u^2(t)dt = \int_0^\infty \left(Ae^{-t/\hat{\tau}} \sin(2\pi\hat{f}_0 t) + Be^{-t/\hat{\tau}} \cos(2\pi\hat{f}_0 t) \right)^2 dt. \quad (3.11)$$

E is proportional to the energy that would be dissipated in the coil of an electrodynamic actuator since the current would be proportional to $u(t)$. Multiplying out the terms we obtain the following expression for the energy:

$$\int_0^\infty A^2 e^{-2t/\hat{\tau}} \sin^2(2\pi\hat{f}_0 t) + B^2 e^{-2t/\hat{\tau}} \cos^2(2\pi\hat{f}_0 t) + 2AB e^{-2t/\hat{\tau}} \sin(2\pi\hat{f}_0 t) \cos(2\pi\hat{f}_0 t) dt. \quad (3.12)$$

However, the last term under the integral can be written as $AB e^{-2t/\hat{\tau}} \sin(4\pi\hat{f}_0 t)$, which time-averages approximately to zero because the decay rate is long compared to the period of oscillation.

$$E \approx \int_0^\infty A^2 e^{-2t/\hat{\tau}} \sin^2(2\pi\hat{f}_0 t) + B^2 e^{-2t/\hat{\tau}} \cos^2(2\pi\hat{f}_0 t) dt. \quad (3.13)$$

$$E \approx \frac{A^2}{2} \int_0^\infty e^{-2t/\hat{\tau}} (1 - \cos(4\pi\hat{f}_0 t)) dt + \frac{B^2}{2} \int_0^\infty e^{-2t/\hat{\tau}} (1 + \cos(4\pi\hat{f}_0 t)) dt \quad (3.14)$$

Equation (3.14), which makes use of the power reduction law from trigonometry, can be simplified again using the fact that the decay rate is long compared to the period of oscillation, so again the sinusoidal terms time-average approximately to zero.

$$E \approx \frac{A^2}{2} \int_0^\infty e^{-2t/\hat{\tau}} dt + \frac{B^2}{2} \int_0^\infty e^{-2t/\hat{\tau}} dt \quad (3.15)$$

Finally solving the integrals we see that the control energy E_P due to changing the resonance frequency and the control energy E_D due to changing the decay time are

approximately orthogonal:

$$E \approx \frac{A^2 \hat{\tau}}{4} + \frac{B^2 \hat{\tau}}{4} = \frac{P_P^2 \hat{\tau}}{4m^2(2\pi \hat{f}_0)^2} + \frac{P_D^2 \hat{\tau}}{4m^2} \approx E_P + E_D. \quad (3.16)$$

Perturbation Analysis

We still need to consider how large P_P and P_D need to be in order to change the resonance frequency and the decay time. The constant $(1 + \eta_P)$ describes the increase in the resonance frequency:

$$2\pi \hat{f}_0 \approx \sqrt{\frac{k + P_P}{m}} \triangleq (1 + \eta_P) \sqrt{\frac{k}{m}} \approx (1 + \eta_P)(2\pi f_0), \quad (3.17)$$

and $(1 + \eta_D)$ describes the decrease in the decay time:

$$\hat{\tau} \triangleq \frac{\tau}{1 + \eta_D}. \quad (3.18)$$

Solving (3.17) and (3.18) enables us to find the appropriate control gains $P_D = R\eta_D$ and $P_P = 2\eta_P k$ assuming that $\eta_P \ll 1$. Making use of (3.16), we can obtain the ratio between the control energies:

$$\frac{E_P}{E_D} = \frac{1}{(2\pi \hat{f}_0)^2} \left(\frac{P_P}{P_D} \right)^2 = \frac{1}{(2\pi \hat{f}_0)^2} \left(\frac{2\eta_P k}{R\eta_D} \right)^2 \approx 4Q^2 \frac{\eta_P^2}{\eta_D^2}. \quad (3.19)$$

We can now set the control energies equal with $E_P \triangleq E_D$ and solve for η_P :

$$\eta_P \approx \frac{\eta_D}{2Q}. \quad (3.20)$$

In other words, if the amount of control power is fixed, the maximum decrease in the decay time is about $2Q$ larger than the maximum increase in resonance frequency! To help put this result in perspective, we provide another equation relating the decay time τ , the fundamental frequency f_0 , and the Q , which holds for $Q \gg 1$ [147]:

$$Q \approx \pi \tau f_0. \quad (3.21)$$

Concrete Example

We consider a concrete example from a low G (G2) note on an acoustic guitar whose frequency is about 100Hz. A perfectly reasonable decay time for the uncontrolled string is $\tau = 2.5$ s. The resonance is musical, so its $Q \gg 1$, and by (3.21), $Q \approx 800$. To decrease the decay time by a factor of 8, we would choose $(1 + \eta_D) = 8$. Applying (3.20), we have

$$\eta_P \approx \frac{\eta_D}{2Q} = \frac{7}{2(800)} \approx 0.0044. \quad (3.22)$$

Hence, using the same amount of control energy, we can only increase the resonance frequency by only

$$\frac{0.0044}{2^{1/12} - 1} \approx \frac{0.0044}{0.0595} \approx 7\text{cents}, \quad (3.23)$$

which is 7/100 of the logarithmic distance in frequency between two adjacent notes. In an approximate sense, this is on the order of the smallest perceptible change in fundamental frequency [59].

Using the same amount of control energy, the decay time can be changed significantly, while the resonance frequency can be changed by only a small amount. This will be a recurring theme throughout this chapter. In general, it is more difficult to change the frequency of a musical resonance than to change its decay time.

3.1.3 Single Resonance Model Predicts General Behavior



It may surprise the reader that we employ the single resonance model of Section 2.2.1 to predict the behavior of collocated control on acoustic musical instruments, which typically have many resonances. Readers with backgrounds in control engineering may be especially surprised since it is well known that the movement of poles can be complicated in general, even for positive real control. In the present section, we explain why the single resonance model is a sufficient control model for predicting the linear control dynamics of positive real musical instruments whose models consist of high Q resonances. It is assumed that the sensor and actuator are collocated.

Example Behavior

We do so by first showing that PID control results in the same behavior as predicted by the single resonance model used in Section 3.1.1. We transform (2.1) into the Laplace domain to represent the transfer function of the musical instrument where $X(s)$ is the displacement and $F(s)$ is the force:

$$\frac{X(s)}{F(s)} = \frac{1}{ms^2 + Rs + k}. \quad (3.24)$$

Now we extend the model to include multiple resonances. For the purpose of creating a suitable, positive real model, we assume that the multiple resonances share the force and that their output velocities sum. Hence, to model N resonances we can sum (3.24) for multiple modes represented by various mechanical parameters.

$$\frac{X_{MULTIPLE}(s)}{F(s)} = \sum_{n=1}^N \frac{1}{m_n s^2 + R_n s + k_n}. \quad (3.25)$$

In the following example, we use $N = 5$ resonances total lining up in a harmonic series. We choose the fundamental frequency $\sqrt{k_1/m_1} = 1000$ radians/s for the convenience of graphically representing the pole positions. Equivalently, the fundamental frequency is about 160Hz. The time constants for the resonances range from 3.3s for the lowest mode to 0.1s for the highest mode. We consider feedback of the form

$$U(s) = \left(P_{DD}s^2 + P_D s + P_P - P_I \frac{a}{s+a} \right) X(s), \quad (3.26)$$

or equivalently considering the velocity $V(s) = sX(s)$ as the sensed variable,

$$U(s) = \left(P_{DD}s + P_D + P_P \frac{1}{s} - P_I \frac{a}{s(s+a)} \right) V(s), \quad (3.27)$$

where $a = 10$ radians/s. The term $a/(s+a)$ functions much like an integrator except that the gain is limited for frequencies beneath 10 radians/s. (3.26) is useful for considering how to implement control given a displacement measurement and to suggest also considering the P_I term. Conversely, (3.27) is more useful for directly

analyzing the positive realness of each control term (see Section 2.3.2). We study the 180° locus for each of the terms independently, assuming that the other gains are set to zero.

$P_P > 0$ Proportional to displacement control

The locus for this case is shown in Figure 3.4 on the left. Since the poles remain near the imaginary axis, their decay times remain long, the imaginary axis can be thought of the “frequency axis” of the poles in radians/s, whereas the real axis can be thought of the “decay time axis.” As is the case for all root loci, as the locus parameter is increased from zero, the closed-loop poles slowly depart from the x ’s, which represent the open-loop poles, and approach the o ’s, which represent the open-loop zeros. Figure 3.4 (left) shows that as P_P is increased from zero, the poles increase in frequency. Since there are two more poles than zeros, the two extra poles approach the asymptotes along the imaginary axis. Conversely, if P_P were to be decreased beneath zero, then the poles would move in the opposite direction, decreasing in frequency (not shown). This behavior is the same as predicted qualitatively by the single resonance model in (3.5).

$P_D > 0$ Derivative of displacement control

The locus for the case $P_D > 0$ is shown in Figure 3.4 on the right. As P_D is increased from zero, the poles move to the left. In other words, $P_D > 0$ decreases the decay time of the poles, and to first approximation for relatively small P_D , the frequency of the poles does not change. Since there is only one more pole than there are zeros, the single extra pole approaches the asymptote to the left along the real axis. Conversely, if P_D were to be decreased beneath zero, the decay time of the poles would increase until at least one pole would become unstable (not shown). This behavior is the same as predicted qualitatively by the single resonance model in (3.4).




Figure 3.4: Root locus for $P_P > 0$ proportional to displacement control on **left** and root locus for $P_D > 0$ derivative of displacement control on **right**

$P_{DD} > 0$ Second derivative of displacement control

The locus for the case $P_{DD} > 0$ is shown in Figure 3.5 on the left. As P_{DD} is increased from zero, the poles decrease in frequency (see Figure 3.5). The decay times of the poles increase slightly, although this effect is minimal and cannot be observed on the scale of Figure 3.5. The number of poles and zeros is equal, so there are no pole trajectory asymptotes. Conversely, if P_{DD} were to be decreased beneath zero, the poles would increase in frequency and their decay times would decrease slightly (not shown). This behavior is the same as predicted qualitatively by the single resonance model in (3.4) and (3.5).

 $P_I > 0$ Integral of displacement control

The locus for the case $P_I > 0$ is shown in Figure 3.5 on the right. Although this controller is technically not positive real, we introduce it here to drive home the advantages of positive real controllers. As P_I is increased from zero, the resonant poles move to the left, causing their decay times to decrease significantly (see Figure 3.5, right). The integral of displacement can in some sense be thought of as a low-pass estimate for the velocity, explaining why this controller behaves like a damper. However, due to its low-pass characteristic, the low frequency poles are more strongly affected by the controller (see Section 3.2.4).

This controller clearly has some utility; however, because it is not positive real, the closed-loop system is not unconditionally stable: the controller becomes unstable if P_I is increased too far. The controller has too many excess poles to be positive real. As a consequence, there are three more poles than zeros on the root locus, so at least one of these poles must enter the right half plane: the pole near the origin due to the leaky integrator moves rightward as P_I is increased and eventually enters the right half plane, destabilizing the closed-loop system (see Figure 3.5, right). Choosing a larger can prevent the pole near the origin from entering the right half plane for small $|P_I|$, but if a is chosen too large, then this pole can adversely affect the fundamental frequency mode.




Figure 3.5: Root locus for $P_{DD} > 0$ second derivative of displacement control on **left** and root locus for $P_I > 0$ integral of displacement control on **right**

In summary, the departure angles of the poles in the cases considered above show that the behavior of PID control for small gains as applied to an acoustic musical instruments with many modes is the same as the effect of PID control when applied to a single musical resonance. We generalize this result for all linear controllers in the following section.

Generalization

Note first that the multi-resonance model given by (3.25) does not directly specify the zeros of the acoustic musical instrument. However, the musical instrument zeros are aligned with the musical instrument poles with respect to the real axis in Figures 3.4 and 3.5. We argue first that this must be the case in general. Consider that we could just as well implement a controller which controls the velocity of the acoustic musical instrument at a point, while the sensing variable would be the force of the acoustic musical instrument on the control actuator. Then the musical instrument plant would become $1/G(s)$. Intuitively, the decay times of the modes would have to be similar because we are controlling an acoustic musical instrument, which is designed to be highly resonant. In other words, when considering impedance control of an acoustic musical instrument instead of admittance control, the poles and the zeros of the plant are reversed; hence, they must have similar decay times. Consequently, the alignment of the poles and zeros of $G(s)$ along the real axis must be similar.

We consider now the departure angle criterion for 180° loci for positive real control. The feedback conforms to the negative feedback convention (see the minus sign in Figure 2.2) since we are implementing positive real control. We wish to investigate the departure angle of the l th open-loop pole p_l . At the l th open-loop pole, we have $G(s)K(s)|_{s=p_l} = -1$. Consequently, $\angle(G(s)K(s))|_{s=p_l} = -\pi$. None of the poles are repeated, so the departure angle criterion for the root locus is that the departure angle ϕ_l of the l th pole is given by

$$\phi_l = \sum_{\forall i} \psi_i - \sum_{\forall i \neq l} \phi_i - \pi, \quad (3.28)$$

where ψ_i is the angle between a line parallel to the real axis and the line segment

from the i th open-loop zero to the l th open-loop pole, and ϕ_i is the angle between a line parallel to the real axis and the line segment from the i th open-loop pole and the l th open-loop pole [65].

The root loci shown in Figures 3.4 and 3.5 are shown for an acoustic musical instrument with $N = 5$ resonances; however, these results generalize for any number of musical resonances. For each resonance added to the instrument, an additional pole and zero must be added to the root loci. Furthermore, since the sensor and actuator are collocated, $G(s)$ is positive real, and hence the poles and zeros must be interlaced, and they cancel each other out when evaluating $G(s)$ at the other open-loop poles. Let \mathcal{O} be the set of indices corresponding to open-loop poles and zeros of resonances other than the l th open-loop pole currently under consideration. Hence, \mathcal{O} does not include any of the poles or zeros due to the controller $K(s)$. We can rewrite (3.28) as follows:


$$\phi_l = \left(\sum_{\forall i \in \mathcal{O}} \psi_i - \sum_{\forall i \in \mathcal{O}} \phi_i \right) + \sum_{\forall i \notin \mathcal{O}} \psi_i - \sum_{\forall i \notin \{\mathcal{O} \cup \{l\}\}} \phi_i - \pi. \quad (3.29)$$

The contribution to ϕ_l due to each pair of angles in the parenthesis cancels approximately, and so we have the result that when calculating the departure angle for the l th pole, we do not need to consider the other musical resonances:

$$\phi_l \approx \sum_{\forall i \notin \mathcal{O}} \psi_i - \sum_{\forall i \notin \{\mathcal{O} \cup \{l\}\}} \phi_i - \pi. \quad (3.30)$$

The essential consequence to be gleaned is that **we can consider each musical resonance independently of the others because the departure angle contributions due to the poles and zeros of the other resonances cancel each other out**. These departure angle contributions are also independent of those due to the poles and zeros of $K(s)$. Hence, for relatively small loop gains, the single resonance model predicts the general behavior of feedback control of an acoustic musical instrument for any linear controller $K(s)$ if the sensor and actuator are collocated.

3.1.4 Controller Phase Response




From the behavior of the previous controllers, we can already draw some conclusions [28]. Consider an acoustic musical instrument with a resonance at frequency f_i .

1. **If at frequency f_i the control signal u lags the velocity \dot{x} by $\pi/2$, then the resonance frequency f_i increases under control. If the magnitude of the feedback signal at f_i is increased, then f_i will increase further.** We know this fact from PID control with $P_{DD} = 0$ and $P_D = 0$. Then $K_{PID}(s) = P_P/s$ with $P_P > 0$, so $\angle K_{PID}(j\omega) = -\pi/2$ for $\omega > 0$.
2. **Similarly, if at frequency f_i the control signal u leads the velocity \dot{x} by $\pi/2$, then the resonance frequency f_i decreases under control.** This behavior will be demonstrated by the “negative spring” in Section 3.2.1. In that case, $P_P < 0$, so $\angle K_{PID}(j\omega) = \pi/2$ for $\omega > 0$.
3. **If the control signal u is precisely in phase with the velocity \dot{x} , then the resonance is damped.** This case corresponds to the PID damping with $P_D > 0$ while $P_P = 0$ and $P_{DD} = 0$.
4. **If at frequency f_i , $\angle K(2\pi j f_i)$ does not match any of the preceding cases, both the resonant frequency and damping of f_i will be changed.**

The above criteria in conjunction with the positive real approach can be used to design many more controllers such as leads and lags, bandpass filters, notch filters, feedforward comb filters, and alternating filters. Many of these controllers are discussed in the next few sections.

3.1.5 Bandpass Control



A bandpass filter can be effectively employed as a controller. It has the advantage that only the frequency region of interest will be affected, so the controller is


less sensitive to nonidealities in other frequency bands. We use the following bandpass filter as a controller because it is positive real:

$$K_{bp}(s) = \frac{\frac{\omega_c s}{Q}}{s^2 + \frac{\omega_c s}{Q} + \omega_c^2}. \quad (3.31)$$

The frequency response of the bandpass filter with center frequency $\omega_c = 2\pi(200\text{Hz})$ and $Q = 20$ is shown in Figure 3.6. The effect of the controller can be determined chiefly by looking at the phase response (see Figure 3.6, bottom). Any instrument resonances located very near ω_c will be damped because $\angle K_{bp}(j\omega_c) = 0$ rad. Instrument resonances located somewhat near to ω_c will be damped, but also shifted some in pitch. Resonances far away from ω_c will be chiefly unaffected by the controller. If we instead invert the sign of the loop gain, then we can selectively apply negative damping at ω_c using this controller.

The mechanical analog of (3.31) is shown in Figure 3.7 in the context of controlling a vibrating string. The existence of the mechanical analog proves that $K_{bp}(s)$ is positive real [161]. Multiple bandpass filters can be placed in parallel in the signal processing chain. They can be tuned independently to affect separate modes concurrently. For instance, a bandpass filter could be tuned to the 1st, 3rd, 5th, etc. harmonics of a vibrating string. Then the controller would make the instrument sound as if it were being played an octave higher. The multiple bandpass filter configuration corresponds to multiple analogs of the type in Figure 3.7 being applied to the instrument such that they are collocated and share the same velocity. Hence, the multiple bandpass filter configuration is also positive real.

3.1.6 Notch Filter Control



A notch filter controller performs the opposite function of the bandpass controller—it applies damping over all frequencies, except for those in the region near




Figure 3.6: Frequency response of the bandpass controller given by (3.31) with $Q = 20$ and $w_c = 2\pi 200\text{Hz}$




Figure 3.7: Mechanical analog of the bandpass controller given by (3.31)




Figure 3.8: Frequency response of the notch filter controller given by (3.32) with $\omega_c = 2\pi(300\text{Hz})$, $Q = 1$, and $\alpha = 10$

the center frequency ω_c . To ensure positive realness, we use the following controller:

$$K_{notch}(s) = \frac{s^2 + \frac{\omega_c s}{\alpha Q} + \omega_c^2}{s^2 + \frac{\omega_c s}{Q} + \omega_c^2}. \quad (3.32)$$

The frequency response is shown in Figure 3.8 for $\omega_c = 2\pi(300\text{Hz})$, $Q = 1$, and $\alpha = 10$. Q is no longer the quality factor as defined in Section 3.1.2, but it still has the same qualitative interpretation. Larger Q 's correspond to “tighter” notches as viewed on the frequency response. The location of the poles and zeros is depicted in Figure 3.9 (right) in the s -plane. Restricting $\alpha > 1$ ensures that the poles lie behind




Figure 3.9: Mechanical analog of the notch controller given by (3.32) (left); Pole-zero plot of (3.32) (right)

the zeros with respect to the imaginary axis of the s -plane (see Figure 3.9).

The notch filter controller can generally be employed to make most of the resonances of the instrument decay quickly, with the exception of any resonances lying significantly far inside the notch. The controller causes somewhat unnatural closed-loop sounding behavior as musicians are not used to musical instruments having a small number of dominant resonances. Note that multiple notches can be created at different frequencies by placing multiple instances of (3.32) in series in the signal processing chain.

3.1.7 Alternating Filter



We now consider how to design a positive real controller that approximates shifting the fundamental frequency of a waveguide, rather than shifting primarily the lowest resonance as with PID control. In other words, what we call the alternating filter should change the pitch while still mostly preserving the harmonicity of the waveguide. In light of the design criteria in Section 3.1.4, we deduce that in order to decrease the fundamental frequency, the controller phase response should lead by $\pi/2$ radians at the waveguide resonances. However, the magnitude response should not decrease at higher frequencies, so a frequency response of the type shown




Figure 3.10: Frequency response of $K_{af}(z)$ where the x 's mark where the transfer function overlaps with the resonances of the waveguide instrument with fundamental frequency $f_0 = 100\text{Hz}$

in Figure 3.10 for the controller $K_{af}(z)$ is needed. The x 's mark the locations of the harmonics for an instrument with fundamental frequency $f_0 = 100\text{Hz}$. Note that the phase response should be nearly constant in the neighborhood of the harmonics so that for small fundamental frequency shifts, the controller phase response is still about $\pi/2$ radians (see Figure 3.10, bottom).

Since the controller is complicated, we design it directly in the digital z -domain rather than first in the s -domain and then applying the bilinear transformation [149]. The definition of z -domain positive realness is similar to the s -domain definition given in Section 2.3.2 [149].




Figure 3.11: Pole zero plot of the digital implementation of $K_{af}(z)$ (zoomed in)

To create the response shown in Figure 3.10, the controller needs to consist predominantly of complex pole and zero pairs placed very near the unit circle as shown in Figure 3.11. Consequently, the controller needs to be implemented using high precision variables. The poles and zeros need to be interlaced in order to satisfy the positive real property [149]. That is, one real pole is placed near DC to start the sequence, and then complex zero and pole pairs are placed along the unit circle as shown in Figure 3.11. The alternating sequence ends with a pole pair, and finally one additional zero is placed at 0.5 to serve qualitatively as a digital lead compensator (see Figure 3.12).

We tested the controller on a digital waveguide simulation of a vibrating string.




Figure 3.12: Pole zero plot of the digital implementation of $K_{af}(z)$

The single sensor and actuator were collocated, and the sampling rate $f_s = 44.1\text{kHz}$ was chosen so that the single sample of controller delay in the simulation would be similar to the controller delay of our laboratory’s digital controller (see Appendix G). A zoomed-out version of $K_{af}(z)$ ’s frequency response is shown in Figure 3.13. The dashed black lines in Figure 3.13 (bottom) placed at the positive real phase boundaries $\pm\pi/2$ reveal that we have sacrificed positive realness of the controller in order to obtain phase lead at higher frequencies. Figure 3.14 shows the frequency response of $K_{af}(z)z^{-1}$, where the single sample of controller delay inherent in the simulation is included. The phase response (see Figure 3.14, bottom) remains within the $\pm\pi/2$ boundaries until comparatively high frequencies. Although the controller becomes nonpassive at these high frequencies, the increased damping in the vibrating string model at high frequencies allows for stable closed-loop dynamics.

We evaluate the simulated performance of the controller by plotting the Fast Fourier Transform (FFT) of the windowed open-loop response in solid blue and the windowed closed-loop response in dash-dotted red (see Figure 3.15). The controller clearly has the effect of decreasing the resonance frequencies of the modes. The modes that have nodes near the location of the modeled sensor and actuator are shifted less than the other modes. According to informal listening tests, we have determined that the closed loop instrument tends to sound more harmonic than the pitch shifting using PID control.

To shift the resonance frequencies more precisely, a more elaborate controller design technique such as pole placement would be required. However, pole placement methods require very accurate models of the instrument. In contrast, the alternating filter controller can be designed to be positive real, making it is less sensitive to musical instrument modeling error and other practical nonidealities. For example, in simulation we found the controller still functioned effectively given sensor and actuator nonlinearity on the order of -20dB THD.




Figure 3.13: Frequency response of the alternating filter $K_{af}(z)$




Figure 3.14: Frequency response $K_{af}(z)z^{-1}$ incorporating the controller delay




Figure 3.15: Fourier transform of windowed open loop response (solid blue) and closed loop response (dash-dotted red)

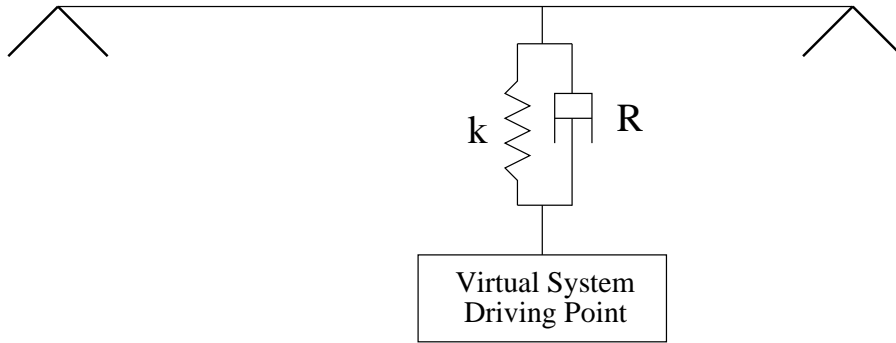



Figure 3.16: Teleoperator-based control of a vibrating string

3.1.8 Teleoperator-Based Control



The previous controllers have the disadvantage there is no direct mapping between the desired closed loop system and controller parameters. However, it would be convenient in musical contexts to be able to specify directly what the dynamic model of the instrument should be [41]. These considerations motivate the development of teleoperator-based control for the feedback control of acoustic musical instruments.


Figure 3.16 shows the mechanical analog for teleoperator-based control of a vibrating string although the idea could be extended to any dissipative acoustic musical instrument. The digital delay does not have a mechanical analog, so it is not shown; however, we assume that the digital control delay is so small that it is negligible. The instrument is linked to the driving point of a virtual dissipative musical instrument by way of the spring k and damper R , which bind the physical and virtual elements together (see Figure 3.16).

The controller depicted in Figure 3.16 is a position-position controller, where the physical instrument is known as the *master* and the virtual instrument is known as the *slave*. Such controllers are frequently used in teleoperation, where a human operator manipulates directly a master robot. A spring-damper controller connects the master robot to a slave robot, which typically completes a task requiring the human's assistance [79].

We have tested this controller in the laboratory using a lightly damped harmonic

oscillator as the virtual instrument. It works sufficiently well; however, the effects of k and R are somewhat deleterious. The spring k detunes resonance frequencies of the virtual model, and the damper R shortens the decay time of the virtual instrument. Nevertheless, the closed loop instrument still behaves qualitatively similarly to the virtual instrument.

3.1.9 Positive Real Multiple-Input Multiple-Output Controllers



Thus far we have assumed that there is only one actuator and one sensor, which are collocated. We have seen that many useful controllers have been represented by mechanical analogs. We noted that increasing the stiffness of the string locally effectively increased the frequency of the lowest partial (see Section 3.1.1). This is physically different than increasing the tension of the entire string in a distributed sense, which directly affects the fundamental frequency [62].

By collocating l sensor/actuator pairs on the acoustic musical instrument, we can obtain a distributed control configuration. Taking $l \rightarrow \infty$ would provide us with unlimited control—we could essentially control the wave equation directly. However, in practice we are limited to a finite number of sensors and actuators. In this section, we consider positive real controllers for collocated multiple-input multiple-output (MIMO) geometries.

We now generalize the positive real formulation to higher dimensional impedances and mobilities. An $l \times l$ real coefficient rational function matrix $K(s)$ of a complex variable s is positive real if the following conditions are satisfied [58]:

1. All elements of $K(s)$ are analytic for $\text{Re}\{s\} > 0$, which is equivalent to stating that the elements of $K(s)$ have no poles in the open right half plane.
2. $K^H(s) + K(s) \succeq 0$ for $\text{Re}\{s\} > 0$, which is equivalent to stating that $K(s)$ is Hermitian positive semidefinite for $\text{Re}\{s\} > 0$.

Condition 2 is equivalent to the following, which is sometimes easier to evaluate:

1. If $j\omega$ is not a pole of any element of $K(s)$, then $K^T(-j\omega) + K(j\omega) \succeq 0$, $\forall \omega \in \mathbf{R}$.




Figure 3.17: Mechanical analog for the controller described by (3.33)

2. If $j\omega_0$ is a pole of an element of $K(s)$, it is a simple pole, and the residue matrix $K_0 = \lim_{s \rightarrow j\omega_0} (s - j\omega_0)K(s)$ in the case ω_0 is finite or $K_\infty = \lim_{\omega \rightarrow \infty} K(j\omega)/j\omega$ in case ω_0 is infinite, is a positive semidefinite Hermitian matrix.

Clearly, we can apply any of the previously discussed positive real controllers to each of the sensor/actuator pairs to obtain a distributed controller represented by a diagonal $K(s)$. However, the MIMO positive real formulation allows new configurations due to the nondiagonal elements of $K(s)$. For instance, mechanical elements can be connected directly between the control points. Figure 3.17 shows one example of a dashpot with parameter P_D and a spring with spring constant P_P connected between two control points. The controller impedance matrix

$$K_{PD}(s) = \frac{P_D s + P_P}{s} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \quad (3.33)$$

is quite simple. However, in simulation we find that these controllers are not particularly useful. For example using $K_{PD}(s)$, the parameter P_P has an effect on the modal frequencies, and the parameter P_D controls damping; however, we can obtain qualitatively similar effects using a single actuator pair and PID control.

MIMO configurations also allow implementing fundamentally different positive real elements such as transformers and gyrators. Multiple instances of these elements can be combined with other positive real elements to synthesize further positive real MIMO controllers. The transformer is less likely to be of practical use in this context because we would need to be able to output a force with one actuator and output a




Figure 3.18: The gyrator element

velocity with the other, so we do not consider the transformer further here. On the other hand, the gyrator, as shown in Figure 3.18, has many useful properties [161]. The manner in which it exchanges the effort and flow variables allows a mass (spring) connected on one of the ports to be seen as a spring (mass) from the other port. Since the gyrator is memoryless, it applies similar amounts of control power to all of the instrument modes, so it will affect them all to similar degrees.

The simplest way to prove the passivity of the gyrator for all $u \in \mathcal{R}$ is to note that the power added to control point 1 $P_{added,1}$ is the same as the power removed from control point 2 $P_{removed,2}$.¹

$$P_{added,1} = F_1 V_1 = (u V_2) \frac{-F_2}{u} = -F_2 V_2 = P_{removed,2} \quad (3.34)$$

Equivalently, passivity of the gyrator can be shown starting from

$$K_{gyrator}(s) = \begin{bmatrix} 0 & -u \\ u & 0 \end{bmatrix} \quad (3.35)$$

and noting that the elements of $K_{gyrator}(s)$ are analytic for all $s \in \mathcal{C}$ and that $K_{gyrator}^H(s) + K_{gyrator}(s) = 0 \succeq 0$ for all $s \in \mathcal{C}$. In fact, the gyrator is the element that makes it possible for $K(s)$ to be Hermitian asymmetric. Driving point mobility and admittance matrices for natural passive linear systems are symmetric because they do not include gyrators [126].

¹In fact, this proof shows that u can be chosen more generally to be any transfer function in order to shape the frequency content of the control.




Figure 3.19: Block diagram for gyrator connecting two independent second-order systems

GYRATOR BETWEEN TWO UNCOUPLED HARMONIC OSCILLATORS



We now wish to consider what would happen if we connected a gyrator between the horizontal and vertical axes of a vibrating string. In reality, there is a small amount of coupling between the axes [149], but for simplicity, we will assume that the coupling is zero. We model each axis as a simple second-order system (see Section 3.1.1). We excite the first system with an additional force $F_{in}(s)$: is not unreasonable for initial analysis.

$$\frac{m_1 s^2 + R_1 s + K_1}{s} V_1(s) = u V_2(s) + F_{in}(s) \quad (3.36)$$

$$\frac{m_2 s^2 + R_2 s + K_2}{s} V_2(s) = -u V_1(s). \quad (3.37)$$

We arrive at the standard feedback control block diagram shown in Figure 3.19.

Since the feedback is negative for all $u \in \mathcal{R}$, no matter what the sign of u , the poles move according to the 180° locus shown in Figure 3.20, and no 0° locus is applicable. Since there are two excess poles, these must approach the two asymptotic trajectories along the $j\omega$ axis as the loop gain approaches infinity. In addition, the departure angles approach 90° and -90° as R_1 and R_2 are taken to be relatively small, which is the case for the dominant modes of musical oscillators. The gyrator causes the resonance frequencies of the modes to be pulled apart (see Figure 3.20). In




Figure 3.20: Root locus for gyrator applied to two independent lightly-damped harmonic oscillators

the multi-mode case of a real vibrating string, where the horizontal and vertical axes are weakly coupled, the harmonics of the vertical and horizontal modes are expected to be driven further away from each other in frequency, resulting in more pronounced beating effects [149].

GYRATOR BETWEEN TWO (COUPLED) POINTS ON SAME AXIS



The effect of connecting a gyrator between two points of a vibrating string in the same axis is considered here. In this case, the modes of the two systems coupled by the gyrator controller are already strongly coupled in open loop, so we must use a more elaborate waveguide model. Let the gyrator (3.35) be applied between the points A and B on a vibrating string as shown in 3.21. The reflection coefficient for velocity traveling waves impinging on the physically almost completely rigid terminations




Figure 3.21: Gyrator connected between the points A and B along a vibrating string

is assumed to be $-p$. Figure 3.21 also shows the labeling of the traveling wave components infinitesimally small distances away from the junctions A and B to the left (L) and right (R). Let τ_A be the wave propagation delay between the leftmost termination and the point A, let τ_B be the wave propagation delay between points A and B, and let τ_C be the wave propagation delay between point B and the rightmost termination.

Let $q = \frac{u}{2R_0}$ to simplify the following. The transfer function between a wave entering point A from the right and the wave exiting A to the right is

$$\frac{V_A^{+R}(f)}{V_A^{-R}(f)} = \frac{-pe^{-j4\pi f\tau_A} + qe^{j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}{1 - qe^{-j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}, \quad (3.38)$$

and the transfer function between a wave entering point B from the right and the wave exiting B to the right is

$$\frac{V_B^{+R}(f)}{V_B^{-R}(f)} = \frac{\frac{-pe^{-j4\pi f\tau_A} + qe^{j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}{1 - qe^{-j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})} (e^{-j2\pi f\tau_B} - q) - q}{e^{j2\pi f\tau_B} + q + q \frac{-pe^{-j4\pi f\tau_A} + qe^{j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}{1 - qe^{-j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}}. \quad (3.39)$$

Taking the limits of (3.38) and (3.39), we obtain the following for the case of infinite control gain, which causes the points A and B to effectively become rigid terminations:

$$\lim_{|u| \rightarrow \infty} \frac{V_A^{+R}(f)}{V_A^{-R}(f)} = \lim_{|u| \rightarrow \infty} \frac{V_B^{+R}(f)}{V_B^{-R}(f)} = -1. \quad (3.40)$$

Furthermore by symmetry, we also have the following, which in conjunction with (3.40) imply that the system becomes three isolated vibrating string segments at infinite control gain:

$$\lim_{|u| \rightarrow \infty} \frac{V_A^{-L}(f)}{V_A^{+L}(f)} = \lim_{|u| \rightarrow \infty} \frac{V_B^{-L}(f)}{V_B^{+L}(f)} = -1. \quad (3.41)$$

Finally, we find the velocity of point B if an external force is exerted on the point B. This is the mobility of the controlled string at junction B be $Y_B(f)$:

$$Y_B(f) = \frac{1}{2R_0} \left(\frac{1 - pe^{-j4\pi f\tau_C}}{1 + p \left(\frac{V_B^{+R}(f)}{V_B^{-R}(f)} \right) e^{-j4\pi f\tau_C}} \right) \left(1 + \frac{V_B^{+R}(f)}{V_B^{-R}(f)} \right). \quad (3.42)$$

We now present an example configuration where $\tau_A = 1.7\text{ms}$, $\tau_B = 1.4\text{ms}$, and $\tau_C = 1\text{ms}$. The open-loop fundamental frequency is $\frac{1}{2} \frac{1}{\tau_A + \tau_B + \tau_C} = 122\text{Hz}$. $p = 0.9$ is chosen artificially small so that the plot is easier to inspect visually; however, it does not alter the qualitative results: **the gyrator causes the resonance frequencies of some of the modes to increase and that of other modes to decrease** (see Figure 3.22). **The decay times of some of the modes increase, while the decay times of other modes decrease.** If the measurements at both points A and B are in phase for a given mode, then the gyrator increases the resonance frequency of that mode. If the measurements at both points A and B are NOT in phase for a given mode, then the gyrator decreases the resonance frequency of that mode. Thus, the effect generally causes the waveguide to become inharmonic. We have verified the behavior qualitatively in simulation using the digital waveguide string model introduced in Section 2.2.2. In some contexts, it might be more desirable to focus the control energy on specific frequency ranges. This could be accomplished by replacing u with a linear filter, represented in the Laplace domain as $U(s)$.






Figure 3.22: Mobility at junction B of the system shown in Figure 3.21

3.2 Positive Real Controllers Motivate The Design Of Nonpassive Controllers




We briefly consider some nonpassive controllers that bear resemblance to some of the positive real controllers that we have developed so far.

3.2.1 Negative Spring Constant



Passive PID control with $P_P > 0$ while $P_D = 0$ and $P_{DD} = 0$ causes the pitch of the controlled instrument to increase (see (3.1)). We now consider PID control with $P_P < 0$, which implements a negative spring. For a given displacement in a direction, the controller exerts a force in the same direction as the displacement. This controller is clearly nonpassive. Consider the controller acting upon a simple mass. If perturbed from the position $x = 0$, the mass would accelerate off to infinity. However, in practice, the controller has utility in controlling acoustic musical instruments since these instruments have some innate stiffness, which cannot be overcome by a negative spring with spring constant sufficiently small in magnitude. By (3.4) and as observed in the laboratory with a vibrating string, $P_P < 0$ causes the pitch of the controlled instrument to decrease.

3.2.2 Negative Damping



We may similarly consider the case of implementing negative damping where $P_D < 0$ while $P_P = 0$ and $P_{DD} = 0$. The controller exerts a force on the instrument in the direction of the velocity, so it is clearly nonpassive. In practice, small amounts of negative damping may often be compensated for by positive damping inherent in the instrument, so by (3.5), negative damping may be employed to increase the decay time of instrument resonances.

3.2.3 Negative Mass



Similarly, a negative mass may be placed at the collocated sensor and actuator on the instrument by choosing $P_{DD} < 0$ while $P_D = 0$ and $P_P = 0$. As long as the net mass is positive, then the resonance's characteristics change according to (3.4) and (3.5).

3.2.4 Integral Control



After investigating the effect of feeding back displacement, velocity, and acceleration in the framework of PID control, we consider feeding back the integral of displacement. However, this controller is unstable, so we use the following similar controller:

$$K_{int}(s) = P_I \frac{a}{s(s + a)}, \quad (3.43)$$

where $a > 0$. If f_0 is the resonance frequency of the lowest open-loop musical instrument mode, then

$$a \ll 2\pi f_0. \quad (3.44)$$

Consequently, the control force is proportional to a filtered quantity that is proportional to the integral of displacement over the audible frequency range. To apply damping with integral control, we choose $P_I < 0$ so that for all frequencies sufficiently far enough away from DC, the controller simulates a sort of “low-pass damper.” That is, for $\omega \gg a$, $\angle K_{int}(j\omega) \approx 0$. Nevertheless, $K_{int}(s)$ is not positive real because the relative degree of $K_{int}(s)$ is two. We investigate the closed-loop stability of this controller applied to the resonance model from Section 2.2.1 with the help of the root locus diagram shown in Figure 3.23. The real axis corresponds to σ and the imaginary axis corresponds to $j\omega$ where $s = j\omega + \sigma$. The pole starting from $-a$ that eventually moves into the right half plane causes the closed-loop system to become unstable for more negative loop gains P_I , for instance if the controller causes the stiffness of the closed-loop system to become negative at DC. Therefore a should be chosen such that strong enough stable control can be obtained for the application while nevertheless




Figure 3.23: Root locus over $P_I < 0$ for $K_{int}(s)$ applied to a lightly damped harmonic oscillator

satisfying (3.44).

Employing a pole perturbation analysis for small $|P_I|$, we obtain a formula for the new decay time constant assuming that (3.44) holds [24]:

$$\hat{\tau} \approx \frac{2m}{R + P_I/(4\pi^2 f_0^2)}. \quad (3.45)$$

Note that (3.45) is essentially the same as (3.4) except that the P_I term adjusting $\hat{\tau}$ rolls off at 12dB/octave in frequency. This is because the control signal is roughly speaking integrated twice in comparison with the control signal for damping via velocity feedback. That is, the integral of displacement is the same as the velocity integrated twice.

3.3 Higher Order Control Schemes



We investigated the performance of a series of controllers in the previous sections of this chapter. We found that changing the decay time of a musical resonance is relatively straightforward, but changing the resonance frequency requires more control power (see Section 3.1.2 in particular). In other words, changing the resonance frequency is difficult. In addition, the controllers we have investigated so far affect the resonance frequencies of modes by different amounts. The alternating filter introduced in Section 3.1.7 has the ability to shift the resonance frequencies of modes by similar amounts, and is in fact a fairly high-order controller. Nonetheless, it detunes the modes of an acoustic musical instrument when shifting them. We shift our focus now to idealized controllers that are capable of changing the resonance frequencies of the modes of a waveguide-based musical instrument without significantly detuning them.

3.3.1 Non-Collocated Traveling Wave-Based Control



Waveguide models, as described in Section 2.2.2, suggest another method for controlling musical instruments. Rather than controlling the physical state, such as the position of the vibrating medium, we may consider controlling the traveling waves propagating within that medium. Consider the linear waveguide string illustrated in Figure 3.24. Let this represent a physical waveguide such as a vibrating string or column of air where the wave variable changes sign at the terminations. For simplicity, we have eliminated the loop filter $H_{lp}(z)$ so that the partials are perfectly harmonic and have the same exponential decay time constant τ [149]²:

$$\tau = -\frac{1}{f_0 \ln(|g|)}. \quad (3.46)$$

The fundamental frequency $f_0 = f_S/(A+B+C+D)$, where f_S is the sampling rate

²If instead one end of the waveguide is inverting and the other non-inverting, for example in the case of a flute, then actually $\tau = -\frac{1}{2f_0 \ln(|g|)}$.




Figure 3.24: Representation of a physical waveguide

in Hz. Since the traveling waves cannot be sensed directly, we must estimate them using arrays of sensors [18][151]. In this section, we consider how sensor nonlinearity prevents this control approach from being implemented for waveguides with long decay time constants. The main point to take away is that the sensor nonlinearity can skew the effective loop gain of the waveguide, which is particularly problematic for long decay time constants. With this meaning in mind, we state that traveling wave-based control is “not robust” to sensor nonlinearity, and we explain further in detail below.

Actuating a traveling wave in a single direction requires either an array of actuators or a single actuator and knowledge of a termination reflection transfer function. Since an analogous statement also holds for sensing a traveling wave, traveling wave-based control cannot be implemented in a perfectly localized fashion at an interior point [18][151]. Nevertheless, to explain this control approach, we initially assume that we can measure the traveling wave exactly in between waveguide portions A and B and that we can actuate the wave variable in between waveguide portions C and D. This structure suggests that we need merely cancel the physical reflection from the right termination and induce an artificial reflection with gain g_{ctrl} as illustrated in Figure 3.25. In this case, the controlled fundamental frequency is $\hat{f}_0 = f_S/(A+D+N)$. Jean Guérard implemented this type of controller for the recorder [73]. He used an array of five microphones placed along the tube to estimate the rightward traveling wave, and then he used a loudspeaker driver at the end of the tube to actuate the leftward traveling wave from the termination (i.e., he chose $C = 0$). We have investigated various methods for estimating traveling wave components based on sensor arrays [18].

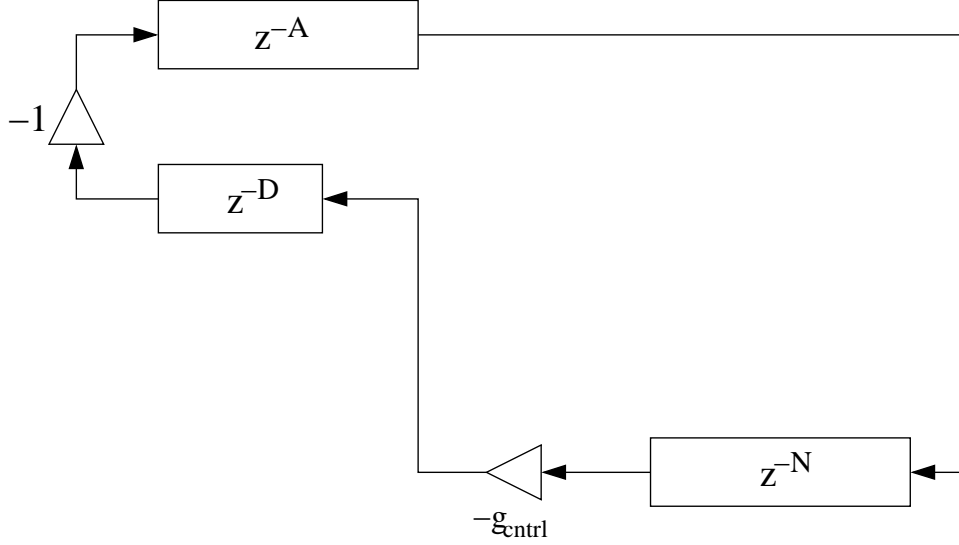


Figure 3.26: Configuration after physical reflection has already been damped.

active or passive damping of the physical reflection, we can argue that traveling wave-based control of a vibrating string cannot be robust in practice given typical vibrating string sensor nonlinearity.

We extend the model to include sensor nonlinearity in Figure 3.27. The coefficient $g_{requested}$ models the loop gain implemented directly by the controller as a standard gain block. Let v_R be the rightward traveling wave component, which is shown exiting z^{-A} and entering z^{-B} in Figure 3.25, and let v_L be the leftward traveling wave component at the same physical location, so as shown in Figure 3.25, it lies within z^{-C} . The physical, measurable velocity $v = v_R + v_L$. \hat{v}_R is the estimate of the rightward traveling wave as corrupted by nonlinear distortion. Hence, $\frac{\partial \hat{v}_R}{\partial v_R}$ models how the linearized loop gain is skewed by the sensor nonlinearity, and this term is incorporated into the model shown in Figure 3.27. To keep the models consistent, we have $g_{ctrl} = g_{requested} \frac{\partial \hat{v}_R}{\partial v_R}$. In essence, $g_{requested}$ is the loop gain we are requesting, but g_{ctrl} is the loop gain we achieve.

$\frac{\partial \hat{v}_R}{\partial v_R}$ depends on the particular sensor nonlinear characteristic curve, on the particular traveling wave estimator employed, and also on the string velocity itself. The exact expression for $\frac{\partial \hat{v}_R}{\partial v_R}$ is especially complicated because the traveling wave estimator has memory. However, we assume that $\frac{\partial \hat{v}_R}{\partial v_R}$ is the same as $\frac{\partial \hat{v}}{\partial v}$ to first order:

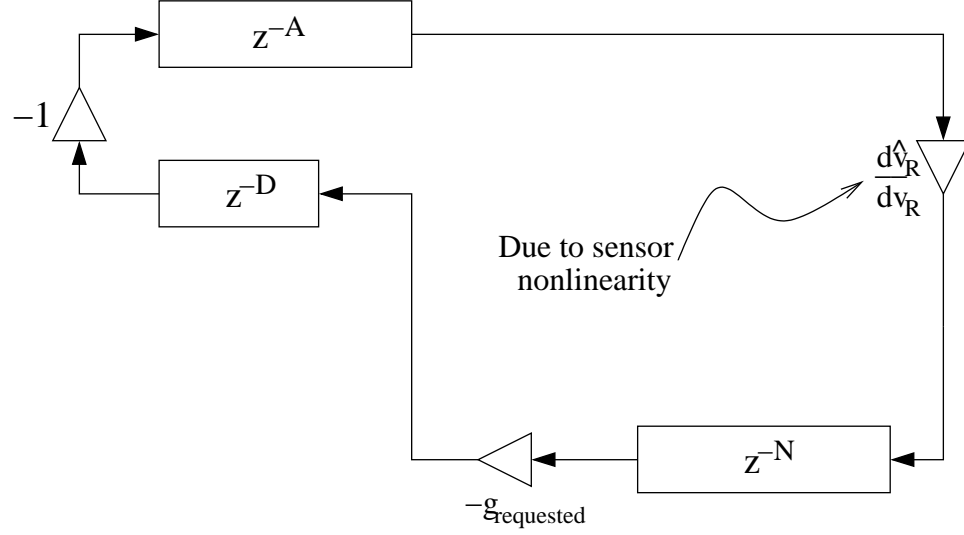


Figure 3.27: Configuration where sensor nonlinearity is modeled by $\frac{\partial \hat{v}_R}{\partial v_R}$

$$\frac{\partial \hat{v}_R}{\partial v_R} \doteq \frac{\partial \hat{v}}{\partial v}. \quad (3.47)$$

This assumption is one of the two reasons why we are making an *argument* in this section rather than a *proof*. This assumption is nonetheless reasonable because we are taking the traveling wave estimator to be linear and time invariant. Hence, if for some particular waveguide state, the estimator happens to compensate for the gain skews of the weakly nonlinear sensors, then we can scale the gain of the state such that the estimator is unable to compensate since superposition does not hold in general for nonlinear systems.

Now we make the argument more specific by choosing a particular sensor nonlinear characteristic. This is the second reason why we are making only an argument and not forming a proof in this section. We now assume that we are measuring the velocity directly, and that the measurement is subject to odd-order harmonic distortion. We further assume that the sensor is operating in a region where the nonlinearity is weak so we approximate the distortion with a single third order term scaled by $\alpha > 0$ in (3.48). \hat{v} is the distorted sensor measurement of the actual velocity v .

$$\hat{v} = v + \alpha v^3 \quad (3.48)$$

The only nonlinearity we are modeling in the entire control system is sensor nonlinearity. We find $\frac{\partial \hat{v}}{\partial v}$ by differentiating with respect to v .

$$\frac{\partial \hat{v}}{\partial v} = 1 + 3\alpha v^2 \quad (3.49)$$

Next we relate α to the total harmonic distortion parameter k (see Section F.1). For all signals v in general, this is not possible, so we choose the signal that is commonly used in measuring harmonic distortion:

$$v = A \cos \omega t, \quad (3.50)$$

where $A > 0$ is the amplitude and ω is the frequency of vibration in radians per second. We can then determine k .

$$\hat{v} = A \cos \omega t + \alpha A^3 \cos^3 \omega t \quad (3.51)$$

$$\hat{v} = A \cos \omega t + \frac{\alpha}{4} A^3 (3 \cos \omega t + \cos 3\omega t) \quad (3.52)$$

$$\hat{v} = (A + \frac{3}{4} \alpha A^3) \cos \omega t + \frac{\alpha}{4} A^3 \cos 3\omega t \quad (3.53)$$

From Equation (3.53), we see that the distortion can either interfere constructively or destructively at ω rad/sec, depending on the sign of α . Given most other nonlinear characteristic functions in place of (3.48), the dependency is much more complicated. We continue by applying the definition of total harmonic distortion (see Section F.1).

$$k = \sqrt{\frac{\frac{1}{2}(\frac{\alpha}{4} A^3)^2}{\frac{1}{2}[(A + \frac{3}{4} \alpha A^3)^2 + (\frac{\alpha}{4} A^3)^2]}} \quad (3.54)$$

Since we are assuming the nonlinearity is weak, we can take $\alpha A^2 \ll 1$ to simplify the expression for k .

$$k = \sqrt{\frac{\frac{1}{2}(\frac{\alpha}{4}A^3)^2}{\frac{1}{2}(A)^2}} = \frac{\alpha}{4}A^2 \quad (3.55)$$

We complete the argument by developing a bound on $|g_{requested}|$ such that the system never enters a state where the linearized loop gain is unstable. That is, we require that

$$|g_{ctrl}| < 1. \quad (3.56)$$

To ensure that $|g_{ctrl}| < 1$, we must have the following:

$$|g_{ctrl}| = |g_{requested}| \left| \frac{\partial \hat{v}_R}{\partial v_R} \right| < 1 \quad (3.57)$$

Then (3.47) and (3.49) imply the following bound:

$$|g_{requested}| \dot{<} \frac{1}{|1 + 3\alpha v^2(x, t)|} \quad (3.58)$$

A sufficient condition for (3.58) can be developed by substituting the maximum value of $v^2(x, t)$ for $v^2(x, t)$.

$$|g_{requested}| \dot{<} \frac{1}{1 + 3\alpha \max\{v^2(x, t)\}} \quad (3.59)$$

Finally, we substitute k into the above inequality by applying (3.50) and then (3.55):

$$|g_{requested}| \dot{<} \frac{1}{1 + 3\alpha A^2} \quad (3.60)$$

$$|g_{requested}| \dot{<} \frac{1}{1 + 12k} \quad (3.61)$$

The value of k that we currently achieve in the laboratory is $k_{velocity}$, as given in Table 3.1. That means that given our laboratory setup, assuming that $\frac{\partial \hat{v}_R}{\partial v_R} \doteq \frac{\partial \hat{v}}{\partial v}$, and assuming that the nonlinearity is of the form (3.48),⁴ we can guarantee that the

⁴Take careful note of how $k_{velocity}$ is determined in Section F.3.1. One additional difference is that the argument assumes we are making a distorted velocity measurement (see 3.48). In contrast, in our laboratory, we are actually making a distorted displacement measurement, which we differentiate

Table 3.1: Relation between distortion and maximum safe waveguide loop gain

k	Value of k	Maximum $g_{requested}$	Decay τ	Description
$k_{velocity}$	-26dB	0.62	3ms	
$k_{displacement}$	-33dB	0.79	6ms	
	-40dB	0.89	14ms	pizzicato pluck
	-60dB	0.988	120ms	normal guitar pluck
	-80dB	0.999	1200ms	very long guitar pluck

linearized loop gain will never become unstable, as long as we choose $|g_{requested}| < 0.62$.

This is a large price to pay to ensure closed-loop stability. Choosing $|g_{requested}| = 0.62$ is equivalent to requesting a system with time constant $\tau = 3\text{ms}$ at $f_0 = 100\text{Hz}$. According to informal listening tests, this is long enough to tell what pitch the note is; however, it is considerably shorter than a typical pizzicato pluck, and it is far shorter than a normal guitar string pluck (see Table 3.1). However, the sensor nonlinearity will cause some partials to decay much more slowly since their effective loop gains will be closer to 1. We also cannot rule out the possibility that some partials will become unstable as a result of nonlinear modal coupling.

We see the same behavior in simulations where sensor nonlinearity is modeled. g_{cntrl} must be chosen small to prevent various harmonics from becoming unstable due to sensor nonlinearity; however, once g_{cntrl} is chosen small enough to ensure stability, the other harmonics decay too quickly to make for a good musical instrument. Clearly such an actively controlled musical instrument would be difficult to use in practice.

Similar arguments to the above could be made for different sensor nonlinearity curves. The above argument suggests that there will be significant stability problems when implementing such a control system, so we believe the non-collocated traveling wave-based approach is not feasible given even minor (e.g. -40dB) sensor nonlinearity as in our laboratory. On the other hand, the method could likely be used for synthesizing tones with fairly long decay times on the order of 0.8 seconds given excellent sensor THD of -60dB or even better (see Table 3.1, line 4).

to obtain a velocity estimate. Strictly speaking, the differentiation and memoryless nonlinearity operators may not have their order reversed in the signal chain, yet we do so in order that this argument remain simple and more easily applicable to other systems.

In contrast, if we allow the traveling wave estimator to be nonlinear, then we can solve the problem trivially in theory. We simply invert the nonlinear characteristic of each sensor and then pass the inverted signals to a linear traveling wave estimator. We decided not to pursue this approach for two practical reasons. On the one hand, each sensor's gain is dependent on conditions such as its temperature. Since a high-power infrared LED is shining directly at the sensor, we assume that the temperature of the sensors will vary over time. To make matters worse, inductive coupling from the actuator corrupts the measurement (see Section F.3.1), further frustrating any sensor nonlinear characteristic inversion.

One may ask how Guérard successfully implemented a non-collocated traveling wave-based controller [73]. First, the microphone sensors he used were likely much more linear than our guitar string sensors. Second, he was not controlling a linear waveguide, but rather a recorder with a saturating nonlinearity at the mouthpiece. The saturating nonlinearity serves to help stabilize the controlled system. Third, recorders have much shorter decay times τ than guitar strings. Nevertheless, Guérard states that some instability problems remained even at the end of his study [73]. It is possible that microphone sensor nonlinearity was one of the factors detracting from the stability of the controlled recorder.

3.3.2 Termination-Based Control



We now consider what we call termination-based control, which involves one sensor and one actuator. This control method is attractive in that it allows specification of any termination reflection function. It allows a physical waveguide to be controlled according to linear digital waveguide virtual instrument synthesis techniques [149].

We remind the reader that higher-order controllers are only necessary for changing the fundamental frequency of a waveguide. If instead the goal is to change the damping of the modes, then velocity feedback (see Section 3.1.1), or bandpass filter control (see Section 3.1.5), or notch filter control (see Section 3.1.6) can be applied using a collocated sensor and actuator pair. If the sensor and actuator are placed near

the termination, then the control system implements a special kind of termination-based control. However, in the following, we consider how to apply termination-based control to change the fundamental frequency of a waveguide.

The sensor and actuator do not need to be collocated, although we will see that collocation has desirable properties in this control application. Consider the configuration for a one dimensional waveguide as shown in Figure 3.28, where the sensor is depicted with a microphone symbol, and the actuator is depicted with a loudspeaker symbol. The controller implements a transfer function with Fourier transform $K(f)$. The ideal sensor and actuator are assumed to operate at points along the waveguide. The impinging rightward-going velocity wave (see Figure 3.28, right) has Fourier transform $R(f)$ at the position an infinitesimally small distance to the left of the actuator. An infinitesimally small distance to the right of the actuator, the rightward-going wave includes the actuated wave $A(f)$ and is thus $R(f) + A(f)$.⁵ $L(f)$ is defined to be the leftward-going velocity wave an infinitesimally small distance to the left of the actuator, so it is directly affected by the actuator. Hence, the leftward-going wave an infinitesimally small distance to the right of the actuator must be $L(f) - A(f)$. $T(f)$ is the reflection transfer function for the physical termination. For an idealized perfectly rigid vibrating string, $T(f)$ would be equal to -1. The propagation delay between the sensor and actuator is $\tau_A = d_A/c$, where d_A is the distance between the sensor and actuator, and c is the wave speed. The distance between the sensor and termination is d_S (see Figure 3.28).

The emerging leftward-going wave $L(f)$ can be expressed as the sum of the actuated wave $A(f)$ and a delayed version of the incoming rightward-going wave and actuated wave as filtered by the termination reflection function.

$$L(f) = A(f) + e^{-j4\pi f(\tau_A + \tau_S)}(R(f) + A(f))T(f) \quad (3.62)$$

$A(f)$ is a function of the controller $K(f)$, the post-actuator rightward-going wave $(R(f) + A(f))$, propagation delays in the system, and the reflection transfer function

⁵This derivation would have been simpler if we had represented the state of the waveguide using force waves rather than velocity waves.

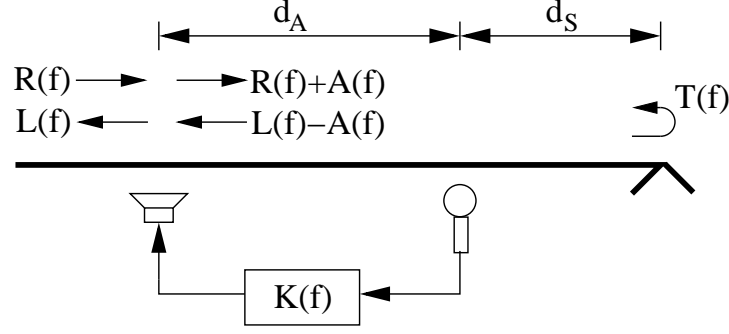


Figure 3.28: Configuration for controlling the termination's transfer reflection function

$T(f)$.

$$A(f) = K(f)(R(f) + A(f))(e^{-j2\pi f\tau_A} + e^{-j2\pi f(2\tau_S + \tau_A)}T(f)) \quad (3.63)$$

Solving (3.63) for $A(f)$ we arrive at the following:

$$A(f) = K(f)R(f)e^{-j2\pi f\tau_A} \frac{1 + e^{-j4\pi f\tau_S}T(f)}{1 - K(f)e^{-j2\pi f\tau_A}(1 + e^{-j4\pi f\tau_S}T(f))}. \quad (3.64)$$

Now we can solve for the induced reflection due to the active termination by substituting (3.64) into (3.62) and solving for $L(f)/R(f)$:

$$\frac{L(f)}{R(f)} = \frac{K(f)e^{-j2\pi f\tau_A}(1 + e^{-j4\pi f\tau_S}T(f))(1 + e^{-j4\pi f(\tau_A + \tau_S)}T(f))}{1 - K(f)e^{-j2\pi f\tau_A}(1 + e^{-j2\pi f2\tau_S})} + e^{-j4\pi f(\tau_A + \tau_S)}T(f). \quad (3.65)$$

The combination of the physical string termination, the sensor, and the actuator can be considered to be an “active” termination. The following derivation is in the spirit of some derivations involving (typically feedforward) control of sound in ducts [115]. However, the derivation differs in that we consider broadband feedback control. The design goal is to choose $K(f)$ such that velocity waves reflect off of the active termination in a desired way. We now solve (3.65) for $K(f)$, and we discover that $K(f)$ can be noncausal if we do not choose $L(f)/R(f)$ carefully.

$$K(f) = e^{j2\pi f\tau_A} \frac{\frac{L(f)}{R(f)} - e^{-j4\pi f(\tau_A+\tau_S)}T(f)}{(1 + \frac{L(f)}{R(f)})(1 + e^{-j4\pi f\tau_S}T(f))} \quad (3.66)$$

We can come up with an easier design method by allowing specification of a controlled reflection transfer function with an inherent delay of τ_A seconds. To this end we define the reflection transfer function $\Lambda(f)$ as follows:

$$\Lambda(f)e^{-j2\pi f\tau_A} \triangleq \frac{L(f)}{R(f)} \quad (3.67)$$

Now when we solve for $K(f)$ we obtain a more convenient expression.

$$K(f) = \frac{\Lambda(f) - e^{-j2\pi f(\tau_A+2\tau_S)}T(f)}{(1 + e^{-j2\pi f\tau_A}\Lambda(f))(1 + e^{-j4\pi f\tau_S}T(f))} \quad (3.68)$$

ACTIVE DAMPING

We first consider how well we can damp vibrations in the waveguide. To have no reflections at all from the active termination, we would need $\Lambda_{damping}(f) = 0$.

$$K_{damping}(f) = \frac{-e^{-j2\pi f(\tau_A+2\tau_S)}T(f)}{1 + e^{-j4\pi f\tau_S}T(f)} \quad (3.69)$$

For a nearly rigidly-terminated string in the laboratory, $T(f) \approx -0.999$ (see Table 3.1), but we see that this will cause the poles of the controller $K(f)$ to be dangerously close to the $j\omega$ axis (see (3.69)). Therefore, **for the controller filter to have well-damped damping poles, we need to have a physical termination that is better damped, for instance $T(f) \approx 0.9$** . This value can be approximately achieved by at least partially passively damping the waveguide termination. This result is analogous to noise-canceling headphone design, where active noise control and passive noise control elements are combined [11].

INDUCING AN INVERTING TERMINATION, PART I

We now consider how to change the fundamental frequency of the waveguide by choosing a constant $\Lambda_{inverting}$ over all frequencies. This technique leaves the closed-loop fundamental frequency a function of the sensor and actuator locations, so it is

not as general as we would like, but the analysis is quite simple, so we begin with this special case. We further assume that the physical termination is already perfectly damped (i.e. $T(f) = 0$). (3.68) then implies the following:

$$K_{inverting}(f) = \frac{\Lambda_{inverting}}{1 + e^{-j2\pi f\tau_A}\Lambda_{inverting}}. \quad (3.70)$$

The poles of $K_{inverting}(f)$ are close to the $j\omega$ axis for $\Lambda(f) \approx -1$, so inducing a perfectly inverting termination is impossible, and inducing a nearly inverting termination will cause $K_{inverting}(f)$ to be poorly conditioned. Consequently, we compromise in the following example by choosing $\Lambda_{inverting}(f) = -0.99$, which is capable of resulting in a fairly short but nevertheless musical decaying note (see Table 3.1). It follows then that $K(0) = -99$. Figure 3.29 shows $|K_{inverting}(f)|$ for the collocated case $\tau_A = 0$ in thick red. **This particular collocated controller simply implements a damper, which is positive real, so we can implement this special case using the velocity feedback term P_D of PID control (see Section 3.1.1).** Figure 3.29 also shows some controller magnitude responses for $\tau_A \in (0, 0.25]$ ms in thin blue lines. These controllers are well-behaved, and they are even positive real—this can be checked by applying the definition of positive real functions in Section 2.3.2 to $-K_{inverting}(s/2\pi j)$. However, these controllers restrict the placement of the artificial termination to a specific position.

INDUCING AN INVERTING TERMINATION, PART II

Now we consider what happens if we want to make it possible to adjust the fundamental frequency to a value which is not as dependent on the sensor/actuator geometry. That is, we now take $\Lambda_{inverting,f0}(f) = \Lambda_{inverting,f0}e^{-j2\pi f\tau}$, where $\Lambda_{inverting,f0} = -0.99$ and $\tau = 0.5ms$. We still retain $T(f) = 0$.

$$K_{inverting,f0}(f) = \frac{\Lambda_{inverting,f0}e^{-j2\pi f\tau}}{1 + e^{-j2\pi f(\tau_A+\tau)}\Lambda_{inverting,f0}} \quad (3.71)$$

The result is plotted in Figure 3.30 for $\tau_A = 0$ in thick red and for $\tau_A \in (0, 0.25]$ ms in thin blue. We see that the gain never exceeds $|K(0)|$, which is about 40dB, so the controller is fairly well conditioned, but the phase lag from the term $e^{-j2\pi f\tau}$ prevents




Figure 3.29: $|K_{inverting}(f)|$ for $\tau_A = 0$ (collocated, thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)

the controller from being positive real. **Thus, when the waveguide fundamental frequency is adjustable over a large range, the controller is not positive real in general.**

INDUCING AN INVERTING TERMINATION, PART III

So far we have assumed quite unrealistically that $T(f) = 0$. Now we consider the case where the physical termination is modestly damped using passive approaches. Now we take $T(f) = -0.9e$ and $\tau_S = 1$ ms. This case reflects the situation in the laboratory better. We retain $\Lambda_{inv,practical}(f) = -0.99e^{-j2\pi f\tau}$, where $\tau = 0.5$ ms. In Figure 3.31, we plot the magnitude response of the controller for $\tau_A = 0$ thick red and for $\tau_A \in (0, 0.25]$ ms in thin blue. The controller's magnitude response is more complicated, but it is still fairly well conditioned because the maximum magnitude is limited. The controller is not positive real due to $\tau \neq 0$.

INDUCING A NON-INVERTING TERMINATION

We consider the case of inducing a **non**-inverting termination. Again for simplicity we assume that the physical termination has already been perfectly damped, so $T(f) = 0$. In this example, we take $\Lambda_{noninvert}(f) = 0.99$. From (3.68), we see now that at low frequencies, the resulting controller operates in the positive feedback mode rather than the negative feedback mode, so it is clearly not positive real. In fact, for $\tau_A = 0$, the controller implements a negative damper. **Inducing a non-inverting**




Figure 3.30: $|K_{inverting,f0}(f)|$ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)




Figure 3.31: $|K_{inv,practical}(f)|$ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)




Figure 3.32: $|K_{noninvert}(f)|$ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)

termination is not positive real.

SUMMARY

Termination-based control can theoretically be applied for creating an active string termination (see Figure 3.28). In the preceding subsections we determined that

1. For the controller filter to have well-damped damping poles, we need the physical termination to be partially damped passively.
2. Inducing a non-inverting termination is not positive real.
3. When inducing an inverting termination using collocated control, the termination-based controller implements a simple damper using PID control.
4. Inducing an inverting termination over a range of positions, and hence making the fundamental frequency adjustable by active control, is not positive real in general.

It is thus not possible to implement positive real control of a one-dimensional waveguide's fundamental frequency using a single wave sensor, single wave actuator, and a single physical termination. Hence, we will not be able to take advantage of the robustness of positive real control (see Section 2.3.2). To make matters worse,

relatively large control gains are required over some frequency regions. In the following section, we show that we cannot achieve large enough loop gains in the laboratory to induce a nearly lossless inverting reflection using the most robust of the termination-based controllers, the damper with damping coefficient much larger than the waveguide's wave impedance. This reasoning explains why we have never tried implementing termination-based controllers in the laboratory, although we believe that some controllers resulting in modestly short decay times could be implemented in practice.

3.3.3 Loop Gain Limitations In The Laboratory



From the previous section, we see that some formulations of termination-based control could be implemented in practice if it were possible to increase the loop gain far enough. In this section, we consider a fundamental special case described in the previous section: the sensor and actuator are collocated, and they implement a damper. Note that this controller is positive real, meaning it should be easier to implement in practice and work better for systems where the natural pole locations are not known.

For small levels of damping, vibrations of the waveguide will be damped, much like how holding a finger against a vibrating string causes it to quickly cease vibrating. However, for larger levels of damping, the controller essentially implements an inverting termination, causing string to split into two more or less independent string segments, each having its own length and fundamental frequency. Given some parameters in practical situations, we estimate how much larger the damping control gain must be to change the fundamental frequency in this manner.

FUNDAMENTAL FREQUENCY

Consider a vibrating string with wave impedance R_0 . Let waves reflect off of the left rigid termination losslessly, and let a dashpot with damping coefficient R_C be attached to the string as shown in Figure 3.33.⁶ Since we are considering velocity

⁶The waveguide could also be from any other one-dimensional waveguide-based instrument, but in this section we assume it is a string so that we can easily consider mechanical analogs of the




Figure 3.33: Dashpot attached to a rigidly-terminated semi-infinite vibrating string




Figure 3.34: Equivalent mechanical configuration to Figure 3.33 where the impedance of the right-hand string segment is lumped into the dashpot

waves, the velocity traveling wave components reflect off of the rigid termination with reflection coefficient -1 (see Figure 3.33). We assume that the dashpot is implemented using feedback control. The dashpot has coefficient large enough to cause a significant reflection with coefficient g off of the dashpot. We do not concern ourselves with what is far to the right of the dashpot, since our goal is to create a rigid termination for the leftmost string segment. Consequently, for convenience we assume that the string continues off to the right-hand side forever supporting zero reflections.

Because we assume no reflections or other waves to arrive on the right-hand segment of the string, we can lump the wave impedance of the string into the dashpot as shown in Figure 3.34. The impedances sum because the string endpoint and the dashpot share the same velocity. For velocity waves, the reflection coefficient g is the

controllers. Other analogs exist for other waveguide-based instruments.

ratio of the impedance step and the impedance sum [149]:

$$g = \frac{R_0 - (R_C + R_0)}{R_0 + (R_C + R_0)} = -\frac{R_C}{R_C + 2R_0}. \quad (3.72)$$

We then solve for R_C .

$$R_C = -\frac{2R_0g}{1+g} \quad (3.73)$$

By substituting in for g using (3.46), we arrive at the following expression:

$$R_{C,change\ f_0} \triangleq R_C = \frac{2R_0 e^{-\frac{1}{f_0\tau_{change\ f_0}}}}{1 - e^{-\frac{1}{f_0\tau_{change\ f_0}}}} \quad (3.74)$$

for $R_{C,change\ f_0}$ as a function of the fundamental frequency f_0 of the segment and the decay time constant $\tau_{change\ f_0}$. In general, the impedance of the dashpot must be much larger than the impedance of the string to cause a strong enough reflection so that the modes with the new fundamental frequency f_0 do not decay too quickly. For instance, for $\tau_{change\ f_0} = 120\text{ms}$ and $f_0 = 100\text{Hz}$, $R_{C,change\ f_0} = 23R_0$. This approximately models the low G2 note on an acoustic guitar. For more information on relating decay times to musical plucked string sounds, see Table 3.1.

DAMPING

To analyze pure damping using a dashpot, we need to consider a right termination as well (see Figure 3.35). g is the reflection coefficient off of the dashpot, so a wave approaching the dashpot continues on to the other segment with gain $(g + 1)$ [91]. Without loss of generality, we assume that the dashpot is closer to the right termination. Now we assume that g is small enough that we can ignore the frequency shifted modes corresponding to the left-hand string segment and right-hand string segment. This assumption will be justified when we show by the end of the argument that $R_{C,damping} \ll R_{C,change\ f_0}$. Thus, we have that velocity waves reflect off of the dashpot and right-hand termination with reflection coefficient

$$g_{eff} = -(g + 1)^2. \quad (3.75)$$




Figure 3.35: Dashpot attached to a rigidly-terminated vibrating string with reflection coefficients labeled for damping analysis

$$g = \sqrt{-g_{eff}} - 1 \quad (3.76)$$

Combining (3.76) and (3.46), we find that

$$R_{C,damped} = -\frac{2R_0g}{1+g} = -\frac{2R_0(\sqrt{-g_{eff}} - 1)}{\sqrt{-g_{eff}}} = -\frac{2R_0(e^{-\frac{1}{2f_0\tau_{damped}}} - 1)}{e^{-\frac{1}{2f_0\tau_{damped}}}}. \quad (3.77)$$



LOOP GAIN COMPARISON

In our laboratory, we can damp vibrations using a dashpot controller so that the decay time of the lowest partials is on the order of $\tau_{damped} = 15\text{ms}$. This is the shortest decay time we can achieve using the current hardware (see Appendix F). This implies that $R_{C,damped} \approx 0.8R_0 \ll 23R_0 = R_{C,change\ f_0}$ for $f_0 = 100\text{Hz}$.

The difference in dashpot damping coefficients actually depends on frequency, so we now compare the ratio of the coefficients $R_{C,change\ f_0}(f_0)$ and $R_{C,damped}(f_0)$ over a range of frequencies. $K_{increase}$, as given below in (3.78), is plotted for $\tau_{damped} = 15\text{ms}$ and $\tau_{change\ f_0} = 120\text{ms}$ in Figure 3.36. $K_{increase}$ becomes larger as f_0 increases. Because we are implementing the dashpot with a digital controller, R_C is proportional to the control loop gain.

As shown in Figure 3.36), it is easier to shift the fundamental frequency for low




Figure 3.36: Increase in loop gain needed for changing the pitch rather than the damping as a function of frequency f_0 ($\tau_{damped} = 15\text{ms}$ and $\tau_{change f_0} = 120\text{ms}$)

f_0 . The cut-off frequency in our laboratory is about 26Hz since $K_{increase}(26\text{Hz}) \approx 1$. However, this limit is clearly a problem since the majority of notes used in most traditional music compositions are higher than this. In fact, the mode at 26Hz would be so low that it would almost be inaudible [59].

$$K_{increase}(f_0) = \frac{R_{C,change f_0}(f_0)}{R_{C,damped}(f_0)} = \frac{\frac{2R_0 e^{-\frac{1}{f_0 \tau_{change f_0}}}}{1 - e^{-\frac{1}{f_0 \tau_{change f_0}}}}}{\frac{2R_0 (e^{-\frac{1}{2f_0 \tau_{damped}}} - 1)}{e^{-\frac{1}{2f_0 \tau_{damped}}}}} \quad (3.78)$$



LIMITATIONS

Now we consider whether we might be able to increase the loop gain far enough to change the pitch in a practical context given the current loop gains we can achieve in the laboratory. Let us assume that for $f_0 \approx 100\text{Hz}$, so we have $K_{increase} \approx 30\text{dB}$.

We have taken efforts to ensure that the *actuator is powerful enough* to provide such large loop gains. The magnets in the actuator are made of neodymium, helping provide relatively large forces (see Section F.2.1). The main limitations here are the

amount of current that can be passed safely along the string and the power provided by the amplifier, but these are not a problem in our laboratory.

In all control systems, the *bandwidth of the control loop* is limited. This characteristic is a consequence of all physical systems being band-limited. As soon as the bandwidth becomes limited, the system evidences phase lag [122]. To ensure that the control system is stable, the magnitude of the gain around the control loop needs to be less than 1 as soon as the bandwidth limitation causes the phase lag to increase further and further (see Section 2.4.1). Consequently, the control loop gain needs to be sufficiently rolled off by the bandwidth limit.

Staying true to our positive real controller philosophy, we are allowed to use only a first-order lowpass filter to roll off the control loop gain. In other words, the controller should behave like a spring at high frequencies. Any higher order lowpass filter in the controller would contribute more than $\pi/2$ radians of phase lag, thus preventing the controller from being positive real. A first-order lowpass filter rolls off the gain by a factor of 10 (20dB) for every increase in frequency of a factor of 10. In our laboratory, a bandwidth limitation causes the “dashpot” to become unstable for loop gains much larger than those needed to achieve $R_C \approx 0.8R_0$.⁷

It follows that $K_{increase}$ further indicates how much smaller we must make the control system bandwidth in order to ensure stability given larger loop gains. Since we currently place the dominant roll-off pole around 4kHz to allow as large loop gains as possible while still damping the audible partials, we would need to place this pole at $4\text{kHz}/K_{increase}$ so that the control system would still be stable after increasing the control gain by $K_{increase}$. This means that without further fundamental improvements to the sensor and actuator design, we cannot induce new fundamental frequencies with a dashpot for partials above approximately $4\text{kHz}/K_{increase} \approx 125\text{Hz}$. This bandwidth is much too small to be musically useful in our application. For example, if the fundamental frequency of vibration $f_0 = 100\text{Hz}$, then the controller would only have the prescribed effect on the lowest resonant frequency. The controller would act like a spring instead of a damper for the higher and clearly audible resonant frequencies.

⁷We believe that inductive coupling between the actuator and sensor may be responsible for this limit (see Section F.4).

We conclude that, in general, it is difficult to obtain large enough loop gains in the laboratory to change the fundamental frequency of a vibrating string using a controller that emulates a damper. It is possible to increase the loop gain by rolling off the damper, turning it into a spring at high frequencies; however, then the bandwidth of the controller becomes limited.

3.3.4 Matched Termination Control



Given the results from Section 3.3.3, we realize we can only hope to change the pitch of the string with control if only relatively small control loop gains are required. In this section, we investigate the possibility of doing so by using a perfectly-matched termination, which improves the conditioning of the control problem. However, we show below that, as in Section 3.3.1, controllers of this form attempting to synthesize long string decay times will suffer from stability problems due to sensor nonlinearity.

String Terminated By Ideal Dashpot

We now assume that the string has a perfectly-matched termination as shown in Figure 3.37. The termination is realized in the form of a damper with the same impedance R_0 as the string's wave impedance. Because the string is perfectly terminated, the termination absorbs all energy from waves impinging on it, reflecting nothing in the absence of control [91]. d is the distance between the sensor and the termination. Since c is defined to be the wave speed on the string, the time delay between the sensor and the termination $\tau = d/c$.

The control system exerts the force $f_F(t) \longleftrightarrow F_F(s)$ at the string termination. Left-going velocity waves induced at the termination are $\frac{f_F(t)}{R_0}$, so the right-going wave $v_R(t)$ at the sensor can be expressed as

$$v_R(t) = v(t) - \frac{f_F(t - \tau)}{R_0}, \quad (3.79)$$

where $v(t)$ is the velocity of the string at the sensor. Accordingly, the controller can




Figure 3.37: Active termination control using a perfectly-terminated string

deduce an estimate $\hat{v}_R(t)$ of the right-going wave at the sensor using the following, where $\hat{v}(t)$ is the estimated velocity of the string at the sensor position:

$$\hat{v}_R(t) = \hat{v}(t) - \frac{1}{R_0} f_F(t - \tau). \quad (3.80)$$

The physical configuration depicted in Figure 3.37 is difficult to implement in the laboratory for many reasons. To name a few, the damper's characteristic constant must be perfectly tuned to the string's wave impedance, the damper must be in contact with the string at an infinitesimally-small point, and the damper must be able to withstand a large horizontal preload force necessary to tension the string.

String Terminated By Servo-Controlled Piston

We present an alternate configuration using a piston-type collocated sensor and actuator as depicted in Figure 3.38. The piston has mass m and spring constant k . An additional damper is placed in parallel with the piston to form the net damping parameter R , making the string termination's damping passively approach that of the string's wave impedance. See Section F.5 for information on this type of motor without R .

Next we close an idealized feedback loop around the piston to ensure that it

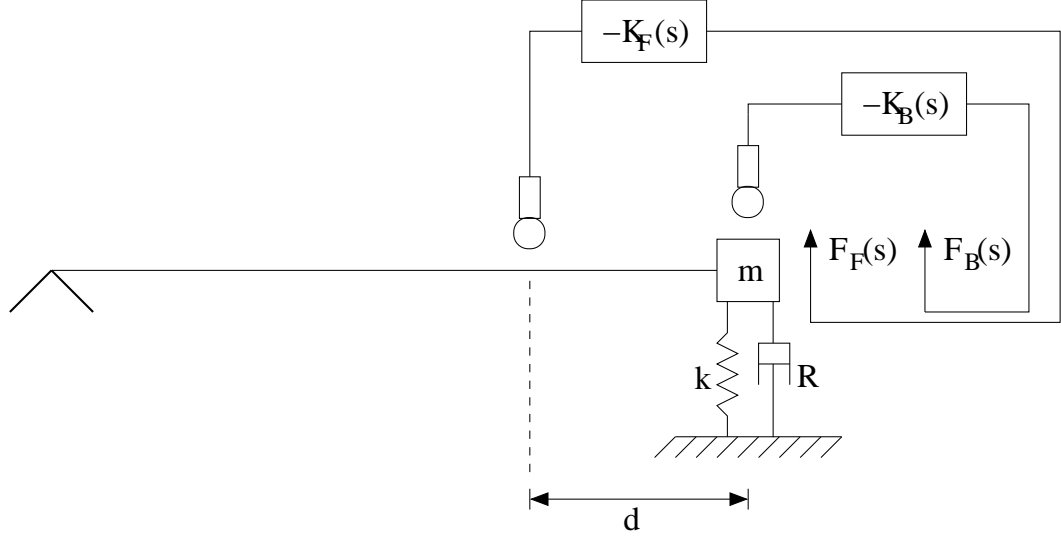


Figure 3.38: Active termination control using a piston-type collocated sensor and actuator

behaves analogously to the system shown in Figure 3.37 at low frequencies for sufficiently small string displacements. The controller cancels out the mass and stiffness of the actuator, and compensates for the damping to result in impedance matching at sufficiently low frequencies. The pole at $-a$ rolls off the mass cancellation at high frequencies.

$$K_B(s) = -ms(a/(s+a)) - k/s + (R_0 - R) \quad (3.81)$$

Effects Of Sensor Nonlinearity

Regardless of how we implement the perfect termination, we have

$$\frac{\partial \hat{v}_R}{\partial v_R} = \frac{\partial \hat{v}}{\partial v} \quad (3.82)$$

due to the form of the traveling wave estimator in (3.80). In other words, the sensor nonlinearity in the velocity estimate $\hat{v}(t)$ corrupts the estimate of the right-going wave $\hat{v}_R(t)$ to the same degree. In this analysis, we can ignore the $\frac{-1}{R_0} f_F(t - \tau)$ term in (3.80) because we know it precisely.

Due to (3.82), we can relate the current controller to the one described in Section

3.3.1. For instance, we may assume the same nonlinearity:

$$\hat{v}(t) = v(t) + \alpha v^3(t). \quad (3.83)$$

We have also perfectly canceled the reflection off of the physical termination, similarly to as in Section 3.3.1, using the servo-controlled piston. Consequently, by the same argument as in Section 3.3.1, we argue the following:

When attempting to provide long decay times, as characteristic for plucked string instruments, practical implementations of active terminations will suffer from stability problems, even for piston-type actuators.

3.4 Summary

The theory of positive real functions can be used for deriving controllers that are guaranteed to result in stable behavior if the acoustic musical instrument is dissipative. Positive real controllers are especially useful for controlling acoustic musical instruments in performance where an exact model of the musical instrument may be unavailable. For instance, the fundamental frequency of vibration may even be unknown in some performance contexts.

We describe a series of positive real controllers in Section 3.1. PID control was previously used in designing feedback controlled acoustic musical instruments [28]. We showed that PID control of the velocity is positive real. We introduced several other positive real controllers to the field of feedback control of acoustic musical instruments: bandpass controllers, notch filter controllers, combinations of bandpass and notch filters controllers, alternating filter controllers, teleoperator-based controllers, and gyrator controllers. We primarily analyzed how they affect only a single musical resonance; however, we showed in Section 3.1.3 that the change in the behavior of a musical resonance is independent of the other musical resonances when subject to linear control.

Positive real controllers are subject to some limitations because they cannot add

energy to the acoustic musical instrument. However, the form of positive real controllers motivates the design of nonpassive controllers. For instance, inverting the phase of a positive real controller by 180° causes the controller to move the poles in the opposite direction. Several controllers designed according to this method were introduced in Section 3.2. These controllers have all been previously applied to the field of the feedback control of acoustic musical instruments except integral control of the displacement, which was analyzed in Section 3.2.4. It can be applied to change the decay times primarily of the musical resonances at lower frequencies.

In Section 3.3, we discussed higher order control methods. We analyzed how they could be used to change the fundamental frequency of one-dimensional waveguide-based instruments. We believe that Termination-Based Control is novel; however, it cannot currently be successfully applied in practice in our laboratory to change the fundamental frequency of a plucked string instrument due to the manner in which control bandwidth limits the maximum loop gain—see Section 3.3.3 for our analysis. In contrast, Non-Collocated Traveling Wave-Based Control has been previously applied to change the fundamental frequency of a recorder [73]; however, as we argued in Section 3.3.1, sensor nonlinearity can destabilize the controller when attempting to synthesize musical resonances with especially long decay times. The same argument prevents Matched Termination Control, which we also believe to be novel, and Termination-Based Control from being applicable to plucked string instruments in our laboratory. In general, the higher order controllers for changing fundamental frequency are aggressive and hence difficult to apply practically to plucked string instruments due to the destabilizing effects of even minor sensor nonlinearity.

The performance of the controllers described in this chapter is further limited by the fact that changing the resonance frequency of musical resonances requires a relatively large amount of control power. Consider controlling a single musical resonance with quality factor Q . If the amount of control power is held constant, the maximum decrease in the decay time of a musical resonance is about $2Q$ times larger than the maximum increase in resonance frequency (see Section 3.1.2). Hence, in these control configurations where we do not have explicit control over the wave speed in the medium, it may not be feasible to change resonance frequencies significantly.

Successful musical instrument designs will likely focus more on adding or removing energy from existing modes. This is why we next investigate controllers that primarily make use of nonlinear strategies for adding energy to existing modes (see Chapter 4) or removing energy from existing modes (see Appendix D). Instrument designers seeking to have significant control over the resonance frequencies of the modes should instead consider directly controlling the wave speed of the acoustic medium.

Chapter 4


Inducing Self-Sustaining Vibrations in Acoustic Musical Instruments

4.1 Introduction

In this chapter, we consider how certain kinds of nonlinear controllers alter the dynamics of acoustic musical instruments in musically appealing ways. In particular, we focus most of our efforts on designing self-sustaining controllers, which can sustain instrument vibrations for an arbitrarily long time. The field of nonlinear control in general is vast, and even simple problems can be hard to solve. To make matters worse, many nonlinear control techniques are difficult to apply if the plant has many modes, such as acoustic musical instruments [138]. Finally, there is a huge divide between nonlinear control theory and how nonlinear control can be applied realistically in practice [98]. For these reasons, we focus on simple controllers that aim to deliver a limited amount of control power to the musical instrument.

Because so many different nonlinear controllers exist, most of the controllers discussed here are motivated by strong nonlinearities present in traditional acoustic musical instruments. Consequently, we suggest that sensors and actuators remain collocated so that passive nonlinearities may be implemented by the controller. For example, this restriction ensures that it is possible to implement nonlinear damping (see Appendix D).

4.1.1 Self-Sustaining Oscillators




In acoustic musical instruments, the most prevalent kind of strong nonlinearity typically gives rise to self-oscillating behavior. At the beginning of play, the energy in the instrument grows until it reaches some equilibrium level. One of the simplest and best known self-sustaining oscillators is the Van der Pol oscillator. The Van der Pol oscillator is not a good model of an acoustic musical instrument because the model order is so low—it essentially models only a single resonance. Nevertheless, it is convenient to start thinking about self-sustaining oscillators in terms of the Van der Pol oscillator because the behavior is comparatively simple. The oscillator's damping parameter $a(x^2 - 1)$ is a function of x [62][162]. For small x , the damping is negative, and for large x the damping is positive.

$$m\ddot{x} + a(x^2 - 1)\dot{x} + Kx = 0. \quad (4.1)$$

Note that we could interpret this system as a physically undamped mass-spring oscillator with mass m and spring constant K , which is augmented with a nonlinear velocity feedback term. The velocity feedback term $a(x^2 - 1)\dot{x}$ drives the mass and spring into a stable oscillation, also known as a limit cycle, at a particular amplitude [138]. In steady state, the damping must be sometimes positive and sometimes negative in order to sustain the oscillation. This modulation causes a strong nonlinear distortion of the velocity feedback.

4.1.2 Sources of strong nonlinearity in acoustic musical instruments are often localized



In contrast with a simple mass-spring oscillator, traditional acoustic musical instruments have many more modes, so they must be modeled using higher order systems. Basic models for many of these instruments consist of a local (often even memoryless) nonlinear element coupled to a higher order musical

resonator.¹ McIntyre, Schumacher, and Woodhouse present a model applicable to the clarinet, the bowed string, and the flute families [108]. For example, a violin string itself is nearly linear, but the interaction between the bow and the string is local and strongly nonlinear. Similarly, a clarinet reed behaves nonlinearly, but it is coupled to a linear waveguide. Rodet and Vergez describe similar models for the clarinet, recorder, bowed string, and trumpet [135].

The dominant nonlinearity in the aforementioned models, with the exception of the trumpet, can be modeled approximately with a nonlinear damper. The important role of the damper should not come as a complete surprise due to its fundamental role in directly relating the force and the velocity.² We consider the role of a nonlinear damper when coupled to a mass and a (nonlinear) spring in Section D.

4.2 Application To Musical Self-Sustaining Oscillators



In this section, we discuss how to design nonlinear controllers reminiscent of particular traditional acoustic musical instruments. Here the controller seeks to induce self-sustaining oscillations in an acoustic musical instrument. In the laboratory, we use a vibrating string to serve as the instrument, although in practice the acoustic musical instrument could be extended to classes of instruments beyond one-dimensional waveguides. For example, it seems likely that it should be possible to bow a cymbal. However, we have not tested instruments beyond the vibrating string, so for now we will make our arguments specific to one-dimensional waveguide instruments. Accordingly, Figure 4.1 reminds the reader of the analogy between controlling a string with a nonlinear damper (Figure 4.1, top) and a vibrating string coupled to a nonlinear damper at some point along the string (Figure 4.1, bottom).

¹Relatively large displacements and velocities can lead to significant distributed nonlinear behavior in some practical situations [81][62]. For example, strings with large displacements behave nonlinearly. Some plates and shells (e.g. cymbals and gongs) also exhibit strong distributed nonlinear behavior even for relatively small displacements [62][158].

²Indeed, in Section 3.3.2, we observed that the damper was also the simplest element that could be used for changing the fundamental frequency of a waveguide.




Figure 4.1: Nonlinear damping control of a vibrating string (top), and string coupled to a nonlinear damper (bottom)

Virtual Bow Characteristic

To bow the instrument with a virtual bow, we need merely choose an appropriate model of a bow-string nonlinear damping curve. The electromagnetically bowed string bears some similarities to a haptic musical instrument (see Chapters 5 - 7), especially in that one can consider that the musician is somehow interacting with a virtual bow; however, the musician is interacting directly with a standard vibrating string element rather than with a haptic device.

Consider the generalized damping characteristic curve shown in Figure 4.2, which is similar to the nonlinear friction that a string experiences when it rubs against a stationary bow. This is an example of nonlinear damping, which is explained in more detail in Appendix D. Instead, to move the implicit virtual bow sideways at a constant velocity v_0 [168] [140], we horizontally offset the characteristic curve to obtain the characteristic shown in Figure 4.3. Part of the characteristic curve now traverses the fourth quadrant, so the damper is no longer passive. That is, whenever $0 \leq \dot{x} \leq v_0$, the damper adds energy to the instrument (i.e. damping is negative). The further v_0 is increased, the larger set of states is where the damping is negative. For many similar characteristic curves and when applied to a waveguide,




Figure 4.2: Damping characteristic curve due to holding a stationary bow held against a string




Figure 4.3: Bowed string nonlinear damping characteristic

the controlled instrument self-oscillates similarly to a bowed string. However, some sets of parameters can lead to period doubling and chaotic behavior [114].

Other Characteristics

To induce sustained oscillations in the instrument according to the sound of other instruments' nonlinearities, we need merely pick the appropriate nonlinear characteristic curve. For example, Figure 4.4 (left) shows the characteristic curve for a clarinet reed [108]. p is defined to be the acoustic pressure and U the volume velocity. p_0 is the mouth pressure control variable. Analogously to the case of the bowed string, when the control variable is zero, the characteristic curve is passive, so all oscillations




Figure 4.4: Clarinet reed nonlinear damping characteristic (left) and trumpet lip nonlinear damping characteristic (right)

will decay. In contrast, as p_0 is increased, the size of the nonpassive state-space grows, generally contributing to an increase in magnitude of self-oscillations.

Figure 4.4 (right) shows the instantaneous nonlinear damping characteristic of a trumpet player’s lips [135]. One major difference now is that the amplitude of the nonlinear characteristic is modulated by the lip position. Nevertheless, the mouth pressure control variable p_0 is still responsible for controlling the size of the nonpassive state-space.

Alternate Velocity Estimate

In practice, and especially when building real musical instruments, a sensor and an actuator might be effectively collocated for only the lowest modes of vibration (see Section 4.7). In this case, nonlinear damper implementations using a velocity estimate may suffer from stability problems at higher frequencies. For this reason, it is worth considering how to make use of a low-pass signal that is approximately in phase with the velocity. Using the methods from Section 3.2.4, we see that we may do so by integrating the displacement x with a slightly leaky integrator and inverting the sign. In general, $\alpha > 0$ for stability (consider the linear case outlined in Section 3.2.4), but $\alpha \ll 2\pi f_0$.




Figure 4.5: Configuration for event-based control of a vibrating string

$$F(t) = R\left(-\alpha \int_0^\infty x(t-\tau)e^{-\alpha\tau}d\tau\right) \quad (4.2)$$

4.3 Event-Based Control

4.3.1 Approach



We investigate another approach known as event-based control, which provides the musician with control over the spectral envelope. Very little theory is available for event-based control [152], so we consider the approach only in the context of controlling vibrating strings, not acoustic musical instruments in general. The approach could be extended to other kinds of acoustic structures, but we have not investigated this aspect. The key elements involved in event-based control are usually an event detector, an observer, and a control signal generator [152]. As in the preceding chapter on linear control, we assume that we do not have a precise model of the instrument we are controlling, so employing an observer is difficult. Instead, we simply measure the velocity at a single position along the string. The configuration with collocated sensor and actuator is shown in Figure 4.5.

In our application of event-based control, we repeat the following indefinitely:

1. Wait until an event is detected.
2. Actuate the system with a predetermined control signal in response.

To reduce the space of the possible controllers, we assume that the control signal is always the same pulse $p(t)$. The goal is to send the pulse into the system periodically to match the unknown or possibly time-varying period of the string. That is, we would like the controller to sustain the vibrations in the string, even though the period of the string may be changing according to what note the musician is playing.

As long as the controller actuates the string with the pulse exactly as the pulse is arriving with the proper sign, the controller will drive the string at its fundamental frequency. This property holds because the string is periodic. We illustrate the property in Figure 4.6, where we assume that the string terminations are memoryless and that wave propagation is nondispersive. These assumptions hold approximately for rigidly terminated vibrating strings with negligible stiffness [149]. Figure 4.6 (a) shows the state of string at rest. Figure 4.6 (b) shows the state of the string velocity directly after the string has been actuated with a pulse. The beginning of the pulse is steeper than the end of the pulse. Blue arrows show the direction of the traveling pulses. Figure 4.6 (c,d,e) show the pulses propagating further over time. The detector is not triggered when the pulses pass under the sensor with inverted sign. Finally Figure 4.6 (f) shows the pulses just before one period has elapsed. They arrive at the collocated sensor in phase. As long as the controller actuates the string with the pulse exactly as the pulse arrives physically, the interference will be constructive. Then the state progresses back to Figure 4.6 (b).

We believe that this control method has not been published elsewhere in the literature. However, Charles Besnainou may have tested a similar method at LAMusicale in Paris, France.³

The dynamic behavior induced when the detector is tracking well bears some resemblance to the slip-stick Helmholtz vibration of a bowed string. When bowing is proceeding according to the Helmholtz string vibration motion, and when the bow is placed near the bridge, the string sticks to the bow during most of the cycle. However, when the Helmholtz corner traveling along the string arrives at the bow, the string begins slipping quickly away from the bow, moving in the opposite direction.

³In personal communication with Adrian Freed, Freed stated that Charles Besnainou had constructed an analog control loop that actuated a vibrating string with a pulse every time that a sensor observed a zero-crossing. Unfortunately, we have not been able to obtain further information.




Figure 4.6: Example of traveling velocity waves on a string actuated with a short pulse

Soon, the Helmholtz corner arrives back at the string approaching from the opposite direction. The Helmholtz corner then causes the string to stick to the bow again [62], and the cycle repeats. The *event* in this scenario is the arrival of the Helmholtz corner.

4.3.2 Bandwidth Advantage

Event-based control offers a unique advantage compared to other practical control techniques described in this thesis. Because of its reliance on event detection, event-based control allows the controller to shape the frequency response of the instrument at relatively high frequencies. For instance, if implemented using a digital controller, event-based control can shape the frequency response up to half the digital sampling rate even if there are anti-imaging filters. In contrast, classical controllers typically require that the magnitude response of the controller rolls off significantly by half the sampling rate in order to retain system stability (see Sections 2.5.3 and 3.3.3).

4.3.3 Analysis



We assume that the string is being actuated with the pulse $p(t)$ at period T seconds, and we analyze the signal $m(t)$ arriving at the listener's ear. We show that the musician can control the spectral envelope of the Fourier transform of $m(t)$ by choosing the shape of $p(t)$.

Let $\mathcal{F}\{g\} = G$ be short hand for the following definition of the Fourier transform:

$$\mathcal{F}\{g(t)\} = G(f) = \int_{-\infty}^{\infty} g(t)e^{-2\pi ift} dt. \quad (4.3)$$

We define the Shah function $III(t)$ to be the infinite pulse train, where $\delta(t)$ is the delta function.

$$III(t) \triangleq \sum_{k=-\infty}^{\infty} \delta(t - k) \quad (4.4)$$

The Fourier transform of the Shah function is the Shah function itself [123].

$$\mathcal{F}\{III(t)\} = III(f) \quad (4.5)$$

The scaling theorem, as given in (4.6), is important because it relates time scaling to frequency scaling.

$$\mathcal{F}\{g(ax)\} = \frac{1}{|a|} G\left(\frac{f}{a}\right) \quad (4.6)$$

We let $v(t)$ be the net signal with which we are actuating the string. It consists of the pulse $p(t)$, repeated every T seconds. This can be represented as a convolution in time with a time-scaled version of the Shah function.

$$v(t) = \sum_{k=-\infty}^{\infty} p(t - kT) = p(t) * III\left(\frac{t}{T}\right) \quad (4.7)$$

If we then define $\mathcal{F}\{v(t)\} = V(f)$ and $\mathcal{F}\{p(t)\} = P(f)$, then we can find $V(f)$ using the scaling theorem (4.6) and the fact that convolution in the time domain corresponds to multiplication in the frequency domain [123].

$$V(f) = P(f) \cdot T \cdot III(fT) \quad (4.8)$$

Finally, there must be some filtering between the velocity induced by the actuator and the sound pressure heard by the ear $M(f)$. Let this filter's transfer function be $H(f)$. The quality of $H(f)$ is primarily due to the resonances of the string and how the string radiates sound; however, of course the filtering between the sound radiated from the instrument to the listener's ear also plays a role. Finally have the following relation:

$$M(f) = H(f)V(f) = H(f)P(f) \cdot T \cdot III(fT). \quad (4.9)$$

If we define f_0 such that $T = \frac{1}{f_0}$, then we have that f_0 is the fundamental frequency of $M(f)$, where $M(f)$ is a sampled version of $H(f)P(f)$ as given in (4.10). $H(f)P(f)$ is the spectral envelope of $M(f)$. The math from this section is illustrated

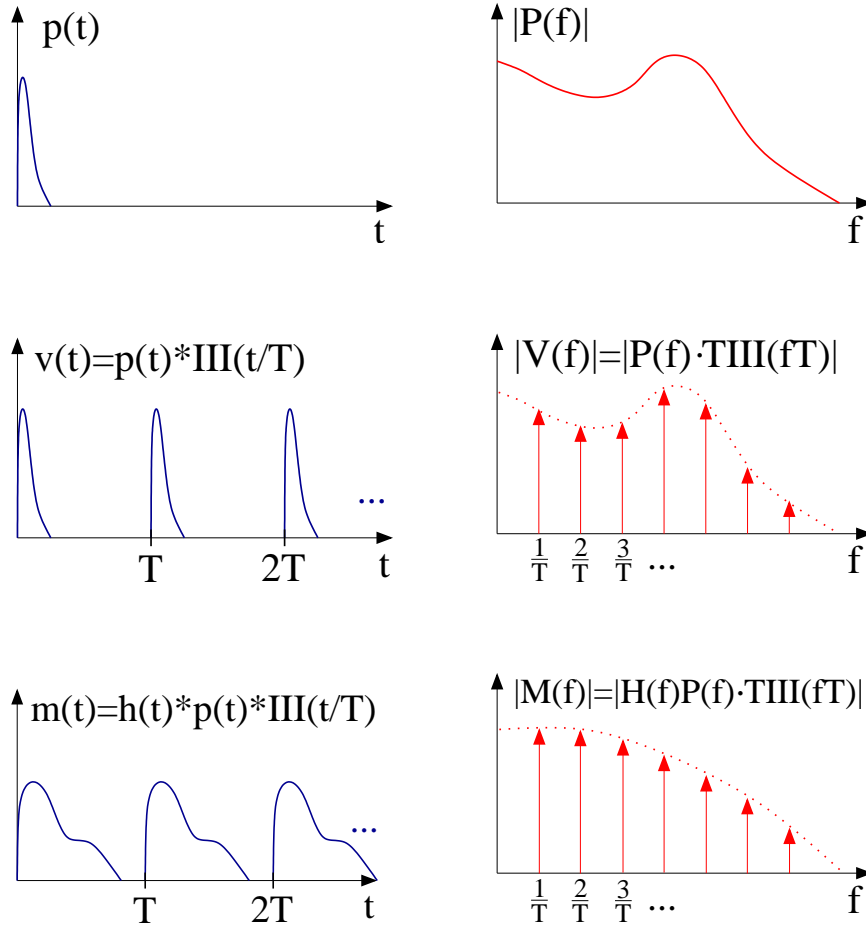


Figure 4.7: Time domain (left) and frequency domain (right) representation of the p , v , and m signals for event-based self-sustaining control of a string. (Note that the illustrated $P(f)$ is not the exact Fourier transform of $p(t)$.)

qualitatively in Figure 4.7.⁴

$$M(f) = H(f)P(f) \cdot \frac{1}{f_0} III\left(\frac{f}{f_0}\right) \quad (4.10)$$

Since we are driving the string at resonance, and the pulses add constructively when the event detector is tracking well, control power is used efficiently. Note that

⁴Figure 4.7 is only qualitative in the sense that $P(f)$ as represented is not the exact Fourier transform of the $p(t)$ shown; rather, $P(f)$ is a hypothetical Fourier transform, which is then sampled by $T \cdot III(fT)$.

the amplitude of $p(t)$ is proportional to the amplitude of $M(f)$, so the amplitude of the $p(t)$ should be chosen according to how loud the controlled instrument should be.

Consider the similarity between how $M(f)$ is formed and how the human voice creates pitched sounds [52]. In the source-filter model, the glottis creates a pulse train of pulses like $p(t)$, and the throat filters these pulses according to some filter that takes on the same role as $h(t)$.⁵

The spectral envelope of a harmonic signal has a strong influence over the perceived timbre of the signal [59]. This controller is especially useful because it allows great freedom in specifying the spectral envelope of $M(f)$. In particular, if $H(f)$ has no zeros, then event-based control can completely specify the spectral envelope of $M(f)$. To achieve the spectral envelope $M_e(f)$, we choose $P(f) = \frac{M_e(f)}{H(f)}$, and then (4.10) implies that $M(f) = M_e(f) \cdot \frac{1}{f_0} III(\frac{f}{f_0})$. It follows that $M_e(f)$ is indeed the spectral envelope of $M(f)$.

Even if $H(f)$ is not known, the mapping from $p(t)$ to $M(f)$ is intuitive for control by the musician. We have implemented this controller in the laboratory for the electric guitar described in Section 4.7. The graphical user interface allows the musician to draw whatever $p(t)$ he or she desires while experiencing how various $p(t)$ cause the controller to induce different self-sustaining vibrations in the electric guitar strings.

4.3.4 Event Detector



The detector observes the string vibration at the sensor and waits for the pulse $p_r(t)$ to arrive. At this time, the string is actuated with the pulse $p(t)$. Finding the optimal event detector is hard because it depends greatly on the dynamics of the string. However, the problem statement implies that we do not even know the period of the string, much less detailed information on the individual modes themselves. In particular, the detector depends on $p_r(t)$, which depends on the pulse $p(t)$ with which we actuate the system. Since we may want to change $p(t)$ over time to vary the spectral envelope, $p_r(t)$ may also change over time. It can be

⁵Note that to match the source-filter model, vowels should be specified via $h(t)$, but we cannot specify $h(t)$ directly using this control method. From (4.10), we see that we may achieve the same result of shaping the spectral envelope by specifying $p(t)$ instead.

hard to implement complex real-time detectors that adjust to $p_r(t)$ changing on the fly. For these reasons, we instead explain one simple, although suboptimal, detector that employs a matched filter.

Let $v(t)$ be the velocity measured by the sensor collocated with the actuator. We then filter the measurement $v(t)$ with the filter having impulse response $z(t)$ to obtain $y(t)$.

$$y(t) = z(t) * v(t) \quad (4.11)$$

We can choose $z(t)$ to optimize our detector. For instance, from matched filter theory, we know that we can maximize the pulse-signal-to-noise ratio of $y(t)$ by letting $z(t) = p_r(-t + \tau)$, where $p_r(t)$ is the pulse we are waiting to receive [93]. τ is chosen so that $z(t)$ is causal but still as short as possible in time. $z(t)$'s simple dependency on $p_r(t)$ allows the detector to easily adapt if $p_r(t)$ changes.

For simplicity, we restrict ourselves to a string in which the velocity waves invert at the terminations. To actuate the string once per period, we should only detect pulses arriving with the same sign as the pulses we are actuating (see Figure 4.6). Our detector, which we have tested in the laboratory on a vibrating string, simply waits for $z(t)$ to exceed a positive threshold. At this time, the string is then actuated with the pulse $p(t)$. If the system were vibrating according to a limit cycle, the detector worked ideally, and the string had memoryless terminations, then the actuated pulse $p(t)$ would perfectly overlap with the arriving pulse $p_r(t)$, and so the string would continue to vibrate perfectly according in a limit cycle. Furthermore, we would then have $p(t) \propto p_r(t)$. However, in the laboratory, the string terminations are not memoryless, and so $p(t) \not\propto p_r(t)$. Nevertheless, for the sake of designing a simple detector, we assume that $p(t) \propto p_r(t)$.


Further heuristic measures improve the detector to achieve suitable behavior despite nonidealities. For instance, in the laboratory, we set a lower limit on the amount of time since a pulse was detected before an arriving pulse can be detected. This limit corresponds to the shortest allowable string period. In practice, the pulses on the string may be smeared due to the detector delay. It may help to make the detector faster by choosing a smaller τ and then truncating $z(t)$ to be causal. Windowing $z(t)$

would indeed improve the pulse-signal-to-noise ratio of $y(t)$, but windowing would cause more delay in the detector, so we do not implement windowing in our detector. We find that our detector works well enough that, within some limits, we can redraw $p(t)$ in real-time without controller tracking problems.

We remind the reader that very little theory is available on event-based control [152]. We have spent some time developing better detectors than the one described above, but they were so complicated that they would have been difficult to implement in real-time in the laboratory. We have chosen to focus our efforts instead on developing other parts of the thesis, but we hope someday we will have the opportunity to look further into developing an optimal detector for event-based control of a vibrating string.

4.4 RMS Level-Tracking Controllers

4.4.1 Dynamic Range Limiter Control



So far we have considered nonlinear controllers that were motivated by sources of nonlinearity in traditional acoustic musical instruments. In this section, we consider another method for developing nonlinear sustaining controllers, and we show that the controlled system is not unstable. The controller continually adjusts the feedback loop gain so that the RMS level of the control signal is approximately constant.

It is important to note that the following controller is not as sensitive to system delay. In fact the stability proof is completely independent of any system delay. This means that the dynamic range limiter controller can be implemented with standard audio hardware, which typically causes system delays longer than the period of the note played on the instrument (see Section G). Composers such as Robert Hamilton and Justin Yang are using controllers of precisely this type. Collin Oldham uses a controller with a limiting nonlinearity to obtain stable behavior despite a system delay longer than the period [26]. The E-Bow and the Sustainiac products also use similar approaches [82][84].

Composers and musicians typically use audio effects to process audio signals. The large class of audio effect signal processors contains too many types of processors to list them all. Many of the common effects include filters, reverberators, amplitude modulators, frequency modulators, distortion generators, and others [149]. Many of the effects are motivated by real physical systems, so in general the effects may be nonlinear and/or time varying. Here we restrict ourselves to effects having only one input and one output. Each of these effects has a particular sound associated with it, and so composers and musicians desire to place these elements in the feedback loops of actively controlled musical instruments. Loosely described, the controlled instrument then generally sounds like some mixture of the uncontrolled instrument and the effect. However, because the various effects behave so differently, it is not straightforward to design a general controller that guarantees some specific behavior for any effect.

We begin with a simple requirement—we wish to prevent the controlled system from becoming unstable. One way to prevent instability is to cascade the effect with a dynamic range limiter, as shown in Figure 4.8 [17]. For example, a simple limiter estimates the RMS level y_p of y using

$$y_p = \sqrt{y^2 * h_{LP}}, \quad (4.12)$$

where $*$ means convolution and h_{LP} is the impulse response of an (often one-pole) low-pass filter with unity gain at DC.⁶ Note that h_{LP} must be non-negative for all time so that the square root in (4.12) is always defined. The output of the limiter u is calculated as follows for some constant P :

$$u = P \frac{y}{y_p}, \quad (4.13)$$

where we assume that y_p is never precisely zero because y always contains some noise

⁶In simulation, we use a slightly more complex RMS level estimate, which makes the limiter less sensitive to noise and small signals. In simulation, we let h_{LP} be a one-pole low-pass filter with unity gain at DC, and we define $y_p = \max(\sqrt{y^2 * h_{LP}}, y_{p,min})$, where the relatively small constant $y_{p,min} > 0$. This change allows the state where the acoustic musical instrument contains no energy to be a stable equilibrium point. Nonetheless, this more complex limiter behavior would make the dynamics analysis needlessly complicated.

in practice. It is generally desirable that y_p changes slowly enough that u does not contain too much harmonic distortion [1] (see (4.43)). However, h_{LP} must allow y_p to change fast enough that the RMS level of u is approximately P .

$$u_p = \sqrt{u^2 * h_{LP}} \approx \frac{P}{y_p} \sqrt{y^2 * h_{LP}} = \frac{P}{y_p} y_p = P \quad (4.14)$$

For the simplicity of the following argument, we assume that the RMS level of u is exactly P . We also assume that the musician cannot add an infinite amount of energy to the acoustic musical instrument, placing the RMS level bound r_{RMS} on the input r . Strictly speaking, the musician can input energy at a different point than the actuator, but for the simplicity of the arguments here, we assume without loss of generality that the musician excites the acoustic musical instrument at the actuator point. This means that the maximum RMS level of the signal arriving at the input to the musical instrument is $r_{RMS} + P$. Finally, let G_{MAX} be the maximum RMS gain of $G(s)$.

$$G_{MAX} = \max_{\omega} G(j\omega) \quad (4.15)$$

G_{MAX} must be bounded because real acoustic musical instruments have some damping at all frequencies [62]. We now form a bound on the RMS level at the output of the instrument x_p :

$$x_p = \sqrt{x^2 * h_{LP}}. \quad (4.16)$$

It follows that the energy in the musical instrument is bounded. This means that the system is not unstable.

$$x_p \leq G_{MAX}(r_{RMS} + P) < \infty \quad (4.17)$$

4.4.2 RMS Level-Tracking Controllers



Here we design a more general RMS level-tracking controller that in many




Figure 4.8: Control configuration with dynamic range limiter in the feedback loop

cases induces a stable limit cycle in the system. The controller should adjust the loop gain L such that x_p is driven to the target x_t . When x_t is held constant, the resulting stable limit cycle may be considered a form of musical sustain. Sustain provides the musician with effective control over the sound because the energy in the system is upper-bounded (assuming observability of all of the modes), yet the instrument is still vibrating, allowing the musician to manipulate the sound in salient ways by adjusting instrument parameters.⁷ In engineering practice, this scenario is unusual because limit cycles are usually undesired. In some sense, the controller is a gain scheduler because it adjusts the loop gain L to drive x_p to x_t ; however, the controller is not a traditional gain scheduler because the instrument's dynamics do not depend on x_p [153].

We cannot develop a controller directly for the system in Figure 4.9 because we have not precisely specified the contents of the *Effect* block. However, by considering only RMS level signals, we may reduce the model order to form the simple approximate model shown in Figure 4.10. The dynamics of the RMS state of the acoustic musical instrument are modeled using a single pole a_p . Since we do not know the contents of *Effect*, we do not know the precise values of a_p and b_p ; however, we at least know the form of all of the system elements. The **max** operation enforces that energy may only be added to the acoustic musical instrument. Technically speaking, the RMS of the instrument state could be reduced by the controller in Figure 4.9,

⁷According to personal communication with Julius O. Smith on Nov. 10, 2009, some unpublished experiments were carried out using a similar method for controlling digital waveguide wind instruments for the Sondius project [127].

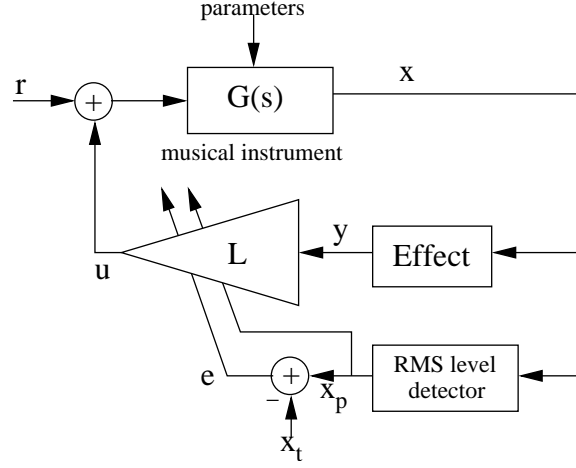


Figure 4.9: RMS level-tracking controller

but this will not be the case in general unless the effect implements damping. The \max operation is in some sense the price we pay for reducing the model order. r_p may be thought of as an excitation corresponding to r . We cause the two controllers to behave analogously, i.e., enforce

$$u_p = Kx_p \quad (4.18)$$

and

$$u_p = \sqrt{u^2 * h_{LP}} = L\sqrt{y^2 * h_{LP}} = Ly_p \quad (4.19)$$

by selecting L according to

$$L = \frac{Kx_p}{y_p}. \quad (4.20)$$

In the following sections, we describe how to complete the controller design, eliminating the final degree of freedom. K is specified for the non-adaptive RMS level-tracking controller in (4.30). L is specified for the adaptive RMS level-tracking controller in (4.49).

Convergence



We now derive the form of RMS level-tracking controllers for which we can bound the rate at which the tracking error $e = x_p - x_t$ converges to zero. At first, we

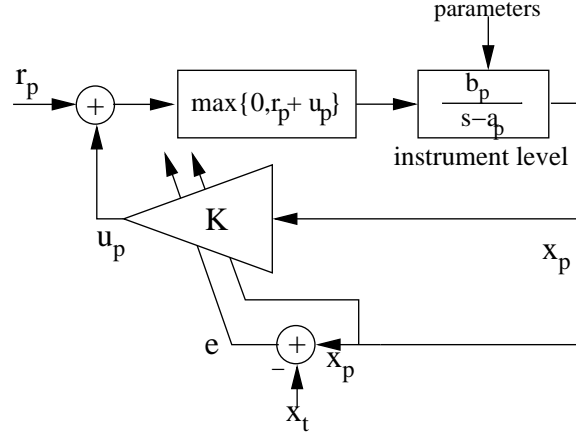


Figure 4.10: RMS level-tracking controller (RMS signals only)

will assume that the musician is not exciting the instrument ($r_p = 0$). The dynamics are

$$\dot{x}_p = a_p x_p + b_p \max\{0, u_p\}, \quad (4.21)$$

where $a_p < 0$. The controller may choose any gain K , and equivalently any u_p , in order to drive x_p to x_t , which is held constant to implement sustain. Consider the Lyapunov function

$$V(e) = e^2/2. \quad (4.22)$$

To provide exponential convergence with time constant less than $1/\lambda$, the controller needs to choose u_p such that

$$\dot{V}(e) < -\lambda V(e) \quad (4.23)$$

for some $\lambda > 0$ [138].

CASE $u_p < 0$: The controller designed in the RMS domain is requesting a negative u_p , which is not possible because u_p is the RMS level of a real signal (see (4.19)). Consequently, we have $\max\{0, u_p\} = 0$, so the control input signal is effectively zero, and the energy in the acoustic instrument must begin decaying. We have $\dot{x}_p = a_p x_p$. Since a_p is stable, both x_p and e will decrease, meaning that we may allow $u_p < 0$ only when $e > 0$. This constraint implies an additional restriction on λ —the controller cannot actively reduce the value of x_p , and so the dynamics of the system must decay

fast enough alone to satisfy (4.23). Consider the following where we are taking $e > 0$:

$$\dot{x}_p = a_p x_p \quad (4.24)$$

$$V(\dot{e}) = e\dot{e} = e(a_p x_p) < -\lambda e^2/2 \quad (4.25)$$

$$a_p x_p < -\lambda e/2 = -\lambda(x_p - x_t)/2 \quad (4.26)$$

$$(a_p + \lambda/2)x_p < \lambda x_t/2. \quad (4.27)$$

We know that $\lambda x_t/2 \geq 0$, so it is sufficient that $\lambda < -2a_p$. In other words, the rate at which we can bound the error e approaching zero is bounded by the speed of the pole $-a_p$.

CASE $u_p \geq 0$: Now we have $\dot{x}_p = a_p x_p + b_p u_p$, which leads to the following:

$$V(\dot{e}) = (x_p - x_t)(\dot{x}_p - \dot{x}_t) = (x_p - x_t)\dot{x}_p = (x_p - x_t)(a_p x_p + b_p u_p). \quad (4.28)$$

By applying (4.23) and after some manipulation, we arrive at the following conditions on the controller for exponential convergence with time constant less than $1/\lambda$:

$$\begin{aligned} u_p &> -\frac{\lambda(x_p - x_t)}{2b_p} - \frac{a_p x_p}{b_p} \quad \text{for } x_p < x_t \\ u_p &< -\frac{\lambda(x_p - x_t)}{2b_p} - \frac{a_p x_p}{b_p} \quad \text{for } x_p > x_t \\ u_p &= -\frac{a_p x_p}{b_p} \quad \text{for } x_p = x_t \end{aligned} \quad (4.29)$$

The light-blue shaded area in Figure 4.11 (left) shows what region valid controller functions must traverse in order to satisfy asymptotic convergence. The dark blue dot in the middle emphasizes the fact that the controller function must include the equilibrium point $(x_p = x_t, u_p = -\frac{a_p x_t}{b_p})$. Qualitatively speaking, the controller must cause x_p to increase if $x_p < x_t$ and cause x_p to decrease if $x_p > x_t$. The dotted line $u_p = -a_p x_p/b_p$ corresponds to the model dynamics. Figure 4.11 (right) shows how the valid controller region is restricted as λ is increased from zero and faster convergence is required.




Figure 4.11: Valid region for convergent controllers (left); Valid region for exponentially-convergent controllers (right)


The system trajectories for $u_p = -a_p x_t / b_p$ are shown in Figure 4.12 (left) [138]. This controller implements the special case of a dynamic range limiter since u_p is held constant [1], so the current work is a generalization of the dynamic range limiter stability result [17]. The light-blue shaded area is drawn to emphasize that some controllers are faster and some are slower (see Figure 4.12 (left)). The dynamic range limiter can be implemented with equivalent controllers according to (4.18) and (4.20):

$$K = \frac{u_p}{x_p} = -\frac{a_p x_t}{b_p x_p} \quad (4.30)$$

$$L = -\frac{a_p x_t}{b_p y_p}. \quad (4.31)$$

We have solved for the gain K needed to control the model shown in Figure 4.10. More practically, we have solved for the gain L to control the *more general model* shown in Figure 4.9. Other more exotic controllers, such as the one whose trajectories are depicted in Figure 4.12 (right), can support multiple equilibria (in dark blue, Figure 4.12 (right)).

Harmonic Distortion



One of the consequences of the RMS level estimation is that L can be made to vary slowly enough that u is only a weakly distorted and scaled version of y . In this section, we estimate the amount of harmonic distortion introduced under favorable




Figure 4.12: Trajectories for the dynamic range limiter controller (left); Trajectories for the controller supporting multiple equilibria (right)

conditions. For the purposes of comparison with the Van der Pol oscillator in section 4.1.1, we assume that the system $G(s)$ describes an undamped mass-spring oscillator. In particular, we have that

$$G(s) = \frac{1}{ms^2 + K}, \quad (4.32)$$

where the frequency of oscillation $\omega = \sqrt{\frac{K}{m}}$ [62]. We further take the low-pass filter $H_{LP}(s)$ in the RMS level detector to be a one-pole filter with unity gain at DC.

$$H_{LP}(s) = \frac{a}{s + a} \quad (4.33)$$

For simplicity we assume that the function $u_p(x_p)$ is constant in the neighborhood of $x_t = x_p$. For example, we could have the special case of the dynamic range limiter where $u_p = \zeta$ for the constant $\zeta \triangleq -\frac{a_p x_t}{b_p}$, which implies the following:

$$u = \frac{\zeta}{x_p} x. \quad (4.34)$$

Next we assume that the system state has converged to the limit cycle, and we estimate the harmonic distortion of the control input u . We assume that the system is vibrating approximately sinusoidally with frequency ω radians/sec, phase offset ϕ_1 ,

and amplitude c_1 , which we will verify at the end of this section.

$$x(t) \approx c_1 \cos(\omega t + \phi_1) \quad (4.35)$$

$$x^2(t) \approx \frac{c_1^2}{2}(1 + \cos(2\omega t + 2\phi_1)) \quad (4.36)$$

$$h_{LP}(t) * x^2(t) \approx \frac{c_1^2}{2} \left(1 + \left| \frac{a}{2j\omega + a} \right| \cos(2\omega t + 2\phi_1 + \angle \frac{a}{2j\omega + a}) \right) \quad (4.37)$$

In order to make use of the Taylor approximation $\frac{1}{\sqrt{1+p}} \approx 1 - \frac{p}{2}$ for $p \ll 1$, we need the low-pass filter to be slow relative to the oscillator's natural frequency.

$$2\omega \gg a \quad (4.38)$$

$$\frac{1}{x_p} = \frac{1}{\sqrt{h_{LP}(t) * x^2(t)}} \approx \frac{\sqrt{2}}{c_1} \left(1 - \frac{1}{2} \left| \frac{a}{2j\omega + a} \right| \cos(2\omega t + 2\phi_1 + \angle \frac{a}{2j\omega + a}) \right) \quad (4.39)$$

$$u = \frac{\zeta}{x_p} x \approx \frac{\zeta \sqrt{2}}{c_1} \left(1 - \frac{1}{2} \left| \frac{a}{2j\omega + a} \right| \cos(2\omega t + 2\phi_1 + \angle \frac{a}{2j\omega + a}) \right) (c_1 \cos(\omega t + \phi_1)) \quad (4.40)$$

$$u \approx \zeta \sqrt{2} \left(\cos(\omega t + \phi_1) - \frac{1}{4} \left| \frac{a}{2j\omega + a} \right| \cos(\omega t + \phi_1 + \angle \frac{a}{2j\omega + a}) - \frac{1}{4} \left| \frac{a}{2j\omega + a} \right| \cos(3\omega t + 3\phi_1 + \angle \frac{a}{2j\omega + a}) \right) \quad (4.41)$$

Finally, since $2\omega \gg a$, we can simplify the approximation (4.41):

$$u \approx \zeta \sqrt{2} \left(\cos(\omega t + \phi_1) - \frac{1}{4} \left| \frac{a}{2j\omega + a} \right| \cos(3\omega t + 3\phi_1 + \angle \frac{a}{2j\omega + a}) \right). \quad (4.42)$$

The approximation, which holds as long as (4.38) is satisfied, for the harmonic

distortion (see Section F.1) introduced by the limiter is


$$k \approx \frac{1}{4} \left| \frac{a}{2j\omega + a} \right|. \quad (4.43)$$

Thus, a parameterizes a trade-off. If a is large then the RMS detector is fast, but there is significant harmonic distortion. The limiting case of a becoming large is similar to placing a clipping nonlinearity in the feedback loop. Electric guitarists use the clipping nonlinearity in guitar amplifier circuits to this end [17]. Some composers also use a clipping nonlinearity instead of a limiter [26].

On the other hand, we can make the harmonic distortion of the control signal u arbitrarily small by making a small, but choosing a too small will make the RMS detector too slow. Choosing $a = 100$ rad/sec reflects a good compromise since this corresponds to fairly fast a time constant of 10ms, while the RMS detector still remains effective for signals containing energy only at frequencies at least a few octaves above $a/2\pi \approx 16\text{Hz}$ [1].

This adjustable trade-off is the primary advantage of RMS level-tracking controllers over a memoryless controller such as the Van der Pol oscillator explained in Section 4.1.1. While we can adjust the amount of harmonic distortion introduced by RMS level-tracking controllers, the harmonic distortion of the control signal regulating the Van der Pol mass-spring oscillator is not easily adjustable—it depends on the parameters m , a , and K and may be large for some sets of parameters [162].

4.4.3 Adaptive RMS Level-Tracking Controllers



We now consider the adaptive scheme shown in Figure 4.13, where the error e in the RMS level of the signal x is used to adapt the loop gain L over time. The controller seeks to drive e to zero.

For a given instrument, effect, and parameters, we may not know what a_p and b_p are. A simple adaptive approach is helpful both in identifying a_p/b_p and allowing for time-varying a_p and b_p due to time-varying instrument and/or effect parameters. Since we are more interested in a general scheme than a fast one, we may start by

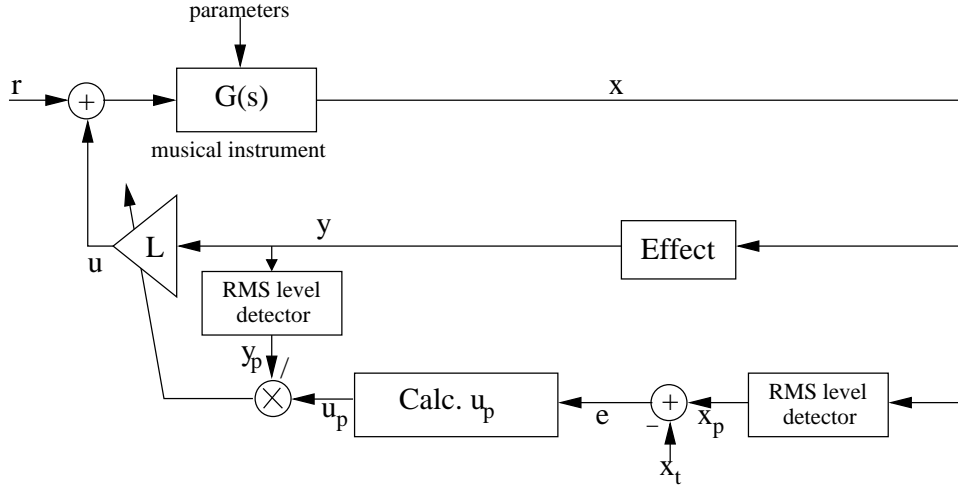


Figure 4.13: Adaptive RMS level-tracking controller

assuming that the controller adaptation is much slower than the plant dynamics, so we have $x_p = -b_p u_p / a_p$. We may now write V as a function of u_p instead of e .

$$V(u_p) = \frac{1}{2}e^2 = \frac{1}{2}(-b_p u_p / a_p - x_t)^2 \quad (4.44)$$

$$\frac{dV(u_p)}{du_p} = e(-b_p / a_p) \quad (4.45)$$

In order to reduce the error over time, we may use the gradient descent method for some $\gamma' > 0$ and $\gamma > 0$.

$$\frac{du_p}{dt} \triangleq -\gamma' \frac{dV(u_p)}{du_p} = -\gamma e \quad (4.46)$$

The condition $u_p < 0$ needs to be avoided because of the **max** operation. One solution in discrete-time implementation is to update the value of u_p only when it results in $u_p \geq 0$. Consequently, we have an RMS signal domain integral controller, which integrates except when further integration would cause u_p to become negative. Let this set of time intervals be defined as “NonNeg.”

$$u_p = -\gamma \int_{\text{NonNeg}} e \, dt \quad (4.47)$$

For instance, to implement (4.47) digitally with a leaky integrator parameterized by $C \approx 1$ where $C < 1$, we could place the following lines of code in the control loop:

```
u_p := C * u_p - gamma*(1-C)*e
if (u_p < 0)
    u_p := 0
```

Adding a proportional term governed by δ can lead to faster convergence and can be thought of as a Proportional-Integral (PI) controller in the RMS signal domain.

$$u_{p,PI} = \max\{0, -\delta e + u_p\} = \max\{0, -\delta e - \gamma \int_{\text{NonNeg}} e dt\} \quad (4.48)$$

Some Wien-Bridge oscillator circuits use this kind of controller to regulate the output amplitude despite any circuit parameter variations [157]. According to (4.18) and (4.20), the loop gain L is then

$$L = \frac{u_{p,PI}}{y_p}. \quad (4.49)$$

If the system adapts slowly enough that u_p never needs to be prevented from becoming negative, then the behavior is linear and the transfer function $X_p(s)/X_t(s)$ may be found.

$$\frac{X_p(s)}{X_t(s)} = \frac{\delta b_p(s + \gamma/\delta)}{s^2 + (-a_p + b_p\delta)s + \gamma b_p} \quad (4.50)$$

Note that the integral term makes the steady-state error go to zero.

$$\left. \frac{X_p(s)}{X_t(s)} \right|_{s=0} = 1 \quad (4.51)$$

Example Including Multiple Equilibria



Next we provide a few examples demonstrating the utility of RMS level-tracking controllers. Simulations were carried out using the digital waveguide model outlined in Section 2.2.2. Here we use a controller similar to the one depicted in Figure 4.12 (right) using the velocity feedback effect where

$$y = \dot{x}. \quad (4.52)$$




Figure 4.14: Ten plucks of linearly increasing magnitude (above), corresponding loop gain $L(t)$ (below)

In practice, there are limits on how fast the instrument may be damped ($L_{min} < 0$) or how fast energy may be added ($L_{max} > 0$), so we would like to restrict $L_{min} \leq L(t) \leq L_{max}$ for all t , which the following accomplishes in addition to inducing multiple equilibria spaced according to P_y :

$$L(t) = L_{max} + (L_{min} - L_{max}) \frac{\text{mod}(y_p(t), P_y)}{P_y}. \quad (4.53)$$

Figure 4.14 shows the results from the simulation where the string is virtually plucked harder and harder over time. Figure 4.14 (top) demonstrates that the final equilibrium RMS state following each pluck depends on the initial plucking condition. Four different stable equilibria in x_p are evident. In comparison with x_p , L is not as smooth due to the modulus operator in (4.53). Figure 4.14 (bottom) reveals that for each pluck, the loop gain L converges to the value 0.08 inducing a stable limit cycle as desired.




Figure 4.15: Driving the RMS of the second harmonic to two values

Example Including Multiple Band-Pass Filters



The band-pass filter effect may also be used advantageously to control the RMS level of a single mode without affecting others. This involves tuning a high-Q band-pass filter h_{BP} (two poles, no zeros) to the mode in question:

$$y = h_{BP} * x. \quad (4.54)$$

If multiple band-pass filters are to be applied simultaneously in order to control separate modes, then Figure 4.9 must be redrawn so that the RMS level of each band-pass *output* $y_{p,n}$ is driven to a target value $y_{t,n}$ as shown in Figure 4.16. Figure 4.15 shows the RMS level of the second harmonic being driven to 130 and then to 25 using the adaptive controller. This simulation provides a good example of how the controller reacts when $r_p \neq 0$. The string is plucked virtually at $t = 0$, which causes the target to be initially overshoot, but the controller is able to adapt to the disturbance by momentarily decreasing L_2 . In this example, eight other harmonics are being controlled simultaneously, although when controlling this many harmonics

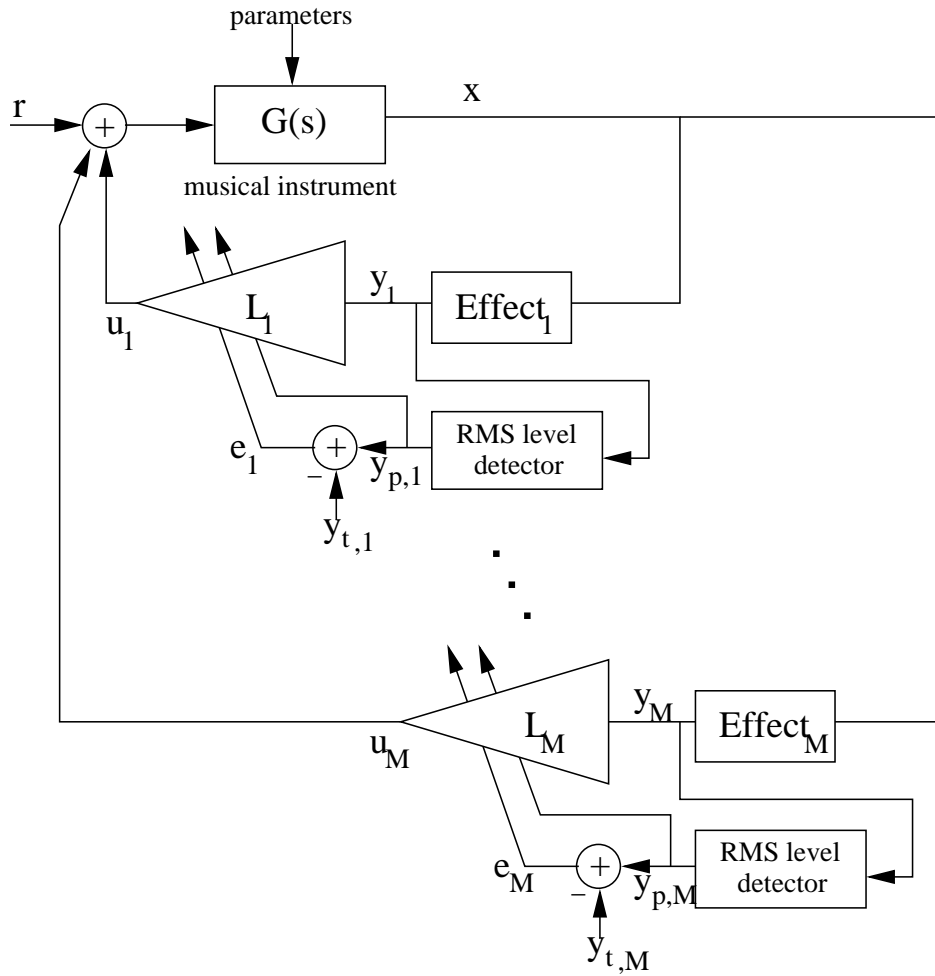


Figure 4.16: System block diagram for implementing multiple RMS level-tracking controllers concurrently

the controller becomes rather sensitive. Nevertheless, the simulations have shown that if a band-pass filter is tuned slightly incorrectly (such as by 2Hz), the limiter can often largely overcome any resulting beating.

Example Including Resonant Ring Modulation



Resonant ring modulation may be implemented using

$$y = x \cdot \cos(2\pi f_c t) \quad (4.55)$$

for some carrier frequency f_c [17]. We will assume that the fundamental frequency of the instrument is f_0 and that the instrument's modes line up perfectly in a harmonic series. Then $f_c = mf_0$ for integer m means that the harmonic series of y will line up with the harmonic series of x . This means that the loop gain L need not be large in order to obtain a limit cycle. On the other hand, if m is not an integer, then the harmonic series do not line up with each other, and a much larger loop gain L is required for obtaining a limit cycle. This strong dependence on effect parameters suggests that the adaptive controller would be useful. The results from a simulation where f_c is swept linearly from f_0 to $2.05f_0$ are shown in Figure 4.17. The dash-dotted red line shows the constant target y_t , and the solid blue line shows the wildly varying RMS level y_p given a fixed loop gain (nonadaptive). In contrast, the dotted green line shows the RMS level y_p when the adaptive controller is used. Besides being able to drive y_p toward y_t without knowledge of the time-varying effective parameter $-b_p/a_p$, the adaptive controller reduces the dynamic range of y_p by 12dB in comparison with the fixed loop gain, nonadaptive controller (compare the solid blue and dotted green lines in Figure 4.17). Such an adaptively controlled instrument is easier for a musician to play because y_p does not depend so much on the instrument and effect parameters.

The companion website provides further explanation of how and why these examples and others were created.⁸ In particular, sound recordings help elucidate why musicians are interested in applying active control to modify acoustic musical instrument behavior.

Examples have demonstrated some of the useful characteristics of both nonadaptive and adaptive RMS level-tracking controllers. Despite low-order modeling and relatively simple controllers, we have induced unusual dynamics in a simulated musical instrument. The controllers are simple and can induce a stable limit cycle for a wide range of audio effects that may be placed in the feedback loop.

⁸<http://ccrma.stanford.edu/~eberdahl/Projects/UnusualDynamics>




Figure 4.17: RMS levels for resonant ring modulation

4.5 Summary

We provide a summary in Figure 4.18 of the characteristics of the nonlinear controllers that we have studied Chapter 4. Passive PID controllers can implement nonlinear damping and stiffness modulation effects. However, in order to induce self-sustaining behavior, we must consider nonpassive nonlinear controllers such as nonpassive nonlinear PID controllers. When these controllers are well-tuned, bowing and vibrating reed behaviors are easily obtained. Event-based control allows direct control over the spectral envelope, but the design of robust event detectors is difficult. Finally, RMS level-tracking controllers can allow the musician to influence the character of the feedback by changing audio effect parameters without affecting the system's stability.

4.6 Conclusions

We have analyzed the most fundamental classes of feedback controllers for changing the acoustics of a vibrating string. Passive controllers are especially useful because the closed loop is stable as long as the instrument is dissipative, and in particular closed loop stability does not depend on the resonance frequency of the instrument modes. Note that passive controllers can be applied to all dissipative acoustic instruments, not only to vibrating strings.





Name	Analysis Class	Linear?	Self-Sustaining?	Passive?	Mechanical Analog	Useful For	Pros	Cons	Implemented?	Simulated?
Passive PID Nonlinear		N	N	Y	Exists	Implementing nonlinear damping and stiffness modulation effects	Stability is guaranteed by passivity. Implements nonlinear modal coupling.	More musically interesting self-sustaining nonlinearities are not passive, so they cannot be implemented.	Y	Y
Nonpassive Nonlinear PID		N	M	N	Exists in some special cases	Inducing self-sustaining behavior using physical processes from traditional musical instruments	Bowing and vibrating reed type behaviors are easily obtainable.	Not all controllers of this class are guaranteed to result in self-sustaining behavior.	Y	Y
Event-Based Control		N	M	N	Does not exist in general	Inducing self-sustaining behavior	The spectral envelope can be directly controlled.	Relatively little theory is available, so designing robust event detectors is difficult.	N	Y
RMS Level-Tracking		N	Y	N	Does not exist in general	Inducing self-sustaining behavior according to metaphors made available by preexisting audio effects	Musicians can usually influence the character of the feedback by changing audio effect parameters without altering the system stability.	No mechanical analogs exist in general for the controllers.	Limiter case	Y

Figure 4.18: Summary of nonlinear controllers covered in Chapter 4

More complex controllers, which aim at changing the termination reflection transfer function of a one-dimensional waveguide, have also been analyzed. While such controllers are indeed interesting, they can require especially large loop gains and can be destabilized by sensor nonlinearity. Hence, they can be applied to instruments such as wind and bowed string instruments where the modes decay relatively quickly; however, their application to plucked stringed instruments is limited because plucked string resonances typically have relatively long decay times.

RMS level-tracking controllers work particularly well in practice. These controllers do not destabilize the closed-loop system because they are designed to satisfy a power-based limitation—the power added to the acoustic musical instrument is upper bounded. Many of these controllers have applications in inducing self-sustaining oscillations in musical instruments. Some of the self-sustaining controllers we have designed are motivated by strong nonlinearities present in traditional musical instruments, while others are novel. In particular, the space of RMS level-tracking controllers provides for a wide range of achievable sounds. Musicians can adjust the sound by placing arbitrary audio effects into the feedback loop, while the controller attempts to adjust the loop gain to induce self-sustaining behavior. Since so many audio effects exist, musicians need to explore the space further to discover new sounds that are achievable with RMS level-tracking controllers.

4.7 Implementation

To help the field continue to progress, we have provided source code for many of our controllers along with an environment for researching, developing, and applying new controllers (see Appendix G). Many of the controllers have been applied to a professional-quality American Fender Stratocaster Deluxe Plus electric guitar. The guitar is shown in Figure 4.19. The piezoelectric sensor at the bridge is made by Graphtech [70], and the actuator is from the Sustainiac electric guitar sustainer system [84]. Robert Hamilton can be seen playing the guitar in Figure 4.20. Please refer to the website on the Feedback Guitar for more details:

<http://ccrma.stanford.edu/~eberdahl/Projects/FBGuitar>



Figure 4.19: Feedback Guitar



Figure 4.20: Robert Hamilton playing the guitar

Chapter 5

Haptic Musical Instruments

5.1 Introduction

So far in this thesis, we have described feedback controlled acoustic musical instruments. In this chapter, we relate them to haptic musical instruments, provide some background on prior research, and motivate our research agenda in employing haptic assistance to alter the gestures that musicians make during performance.

5.2 Overview

In the previous chapters, we described and introduced feedback controlled acoustic musical instruments, and we encountered some limitations. For example, significant power is required to change the resonant frequency of a physical musical resonance (see Section 3.1.2). In addition, since the wavelengths of audible sound vibrations can be short, sensors, actuators, and control system delay further limit the bandwidth over which an acoustic musical instrument can be feedback controlled in practice (see Section 3.3.3). In order to reduce the effects of these limitations, more elaborate designs are required, which tend to be less convenient. For example, extra effort is required to make the designs portable. Consequently, some challenges remain in deploying feedback controlled acoustic musical instruments widely.

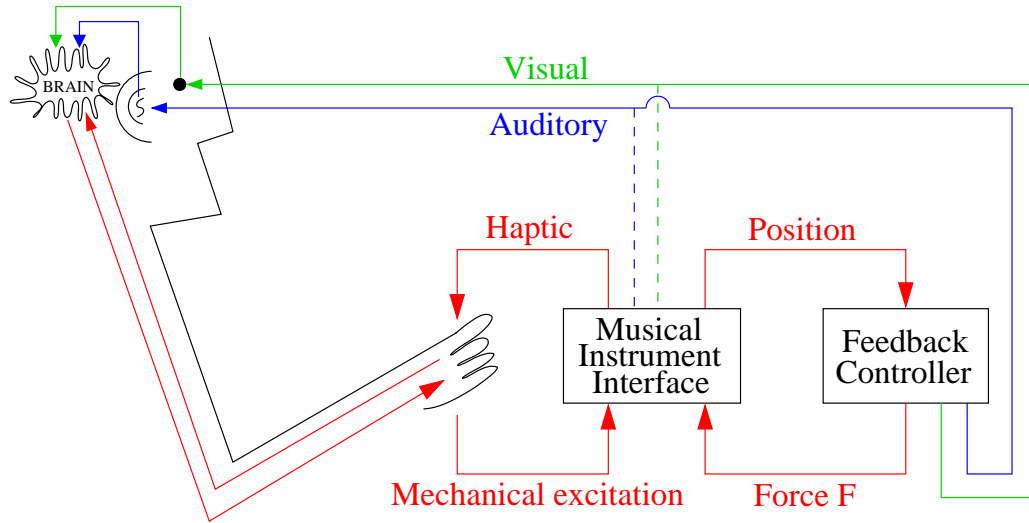


Figure 5.1: Musician interacting with a typical haptic musical instrument

Furthermore, in practical musical contexts, humans are incapable of feeling vibrations at frequencies higher than 500Hz [30]. This psychophysical perceptual limit suggests that no feedback controlled acoustic musical instrument needs to present hard-to-control haptic vibrations at high frequencies to musicians. Instead, musical instrument designers can simulate the dynamics of musical instruments using digital feedback controllers while allowing the musician to interact mechanically with the instrument at vibrations beneath 500Hz. Consequently, in the remainder of this thesis, we focus on designing what we call *haptic musical instruments*. These instruments can currently be employed more easily in live performance contexts due to the ease of control and the availability of commercial haptic devices.

A haptic musical instrument (HMI) is an electronic musical instrument that provides the musician not only with auditory feedback but also with haptic (or force) feedback. Although already introduced in Section 1.4, we present Figure 5.1 again to remind the reader of the representation of a musician interacting with a typical haptic musical instrument. The instrument consists of a musical instrument interface, which could for example be a robotic arm, whose dynamics are controlled by a feedback controller. The primary forms of visual and auditory feedback are provided by the feedback controller, generally by way of a graphical display and a loudspeaker. The

connection from the musical instrument interface to the visual and auditory feedback paths are shown dashed (see Figure 5.1) since, depending on the application, they may play a less important role in affecting a musician's interaction with a haptic musical instrument.

5.3 Controlling Acoustics

Most haptics research and development employs haptic systems to model virtual reality. In this application, the feedback controller is programmed to simulate a virtual environment. For example, the controller could simulate surgery to help surgeons practice surgery [112], or the controller could simulate a remote environment allowing a surgeon to carry out surgery at a distance [74].

In musical contexts, the virtual environment can enable a musician to interact with a virtual musical instrument. To this end, the feedback controller need merely emulate the physics of the musical instrument. The auditory feedback corresponds to the vibrations of the virtual instrument, and any visual feedback can provide the musician with more information on the musical interaction. Due to the advantages of passive control (see Section 2.3.2), we suggest that the virtual instrument be connected to the musical instrument interface according to the positive real philosophy. In other words, the virtual musical instrument model should include a port at the point of haptic interaction, the sensors and actuators of the musical instrument interface should be collocated, and the interface and port should be connected together by a passive network, such as a spring. Then the control system is guaranteed to be stable [2]. Musical instrument designers may wish to consult Chapter 2 for more details on how to control acoustics according to the positive real philosophy.

Although not specifically stressed in their publications, researchers who previously designed haptic musical instruments that mimic real acoustic musical instruments, have implicitly followed the positive real philosophy by utilizing collocated control. See Section A.1 for a literature review of these instruments. However, the physical models these researchers employed have suffered from efficiency, accuracy, and

calibration limitations. Appendix A presents a novel method for allowing for haptic-acoustic interaction with a virtual musical instrument, which is implemented using digital waveguides and overcomes the limitations of prior models for virtual haptic musical instruments. Examples are provided of a musician haptically bowing and plucking a virtual vibrating string (see Appendix A).

5.4 Controlling Gestures

Up until this point, we have studied the feedback control of a musical instrument's acoustics. However, feedback controllers are capable of doing much more than emulating the physics of acoustic musical instruments. In fact, the programmability of feedback controllers endows them with the capability of modifying the very nature of the way in which the musician interacts with an instrument. From now on, we will study the feedback control of the musician's *gestures* to a musical instrument.

5.4.1 Motivation

We frame this research application in the context of prior research by considering the paradigms of musical practice represented in the two axis space shown in Figure 5.2. The horizontal axis represents to what degree the practice can be used for expressive performance in live settings. We consider the opposite of live expressiveness to represent practice in which all parameters are precisely controlled and prespecified by the composer. The vertical axis represents how free the practice is from physical constraints versus how physical any interaction with the musician may be. Of course many other axes are important, but we consider only those that are most relevant to the discussion here. The interested reader may enjoy reading about the seven-axis space for comparing musical interfaces introduced by Birnbaum et al. [31].¹

The most traditional music making practice is represented by the woman shown playing the cello in Figure 5.2. She controls numerous variables simultaneously, allowing her to play expressively in a live scenario. We chose to use a bowed string

¹Miranda and Wanderley introduce another scheme in their book in which musical interfaces are classified into four groups [110].




Figure 5.2: Paradigms for musical practice

instrument in this representation because of the history of expressive performance with bowed stringed instruments [63]. Besides traditional performance interactions, the musician could also engage in any number of extended interactions involving scraping the acoustic instrument, burning it, or throwing it, etc. While these interactions would likely damage the acoustic instrument, they would nevertheless be expressive and result in some physical sounds. The interaction is only as expressive as the musician makes it, and the musician relies on the instrument for supporting expressive interaction. In this traditional paradigm, the interaction between the musician and the music-making device is also entirely physical, which explains why the woman playing the cello is placed in the bottom right of Figure 5.2.

Perhaps the most diametrically opposed music making practice corresponds to a computer autonomously running a program to generate a sound signal, which is fed to a loudspeaker. The representation of this practice is shown in the upper left corner

of Figure 5.2 since the sound can be precisely controlled by programming. This practice could also be represented by a digital music player, which upon a button press, plays prerecorded music and supports only the most minimal live physical interaction. It is nearly free from physical constraints—in other words, very little physical interaction is involved. As a consequence, the timing and pitches of music generated using this paradigm can be so precisely predetermined, that such music is described by some critics as sounding “synthetic,” “mechanical,” or “dry” because no human instrumentalist participates directly [29][170].

Another important practice paradigm involves a musician playing a digital keyboard-based instrument or other digital interface. This practice can be quite expressive in live contexts, where a musician can press a key immediately bringing about a result. Additional sensors such as key velocity sensors, accelerometers, etc., can be used to provide the musician with additional channels for controlling the sound. The practice is freed from some physical constraints as the sound can be calculated by a computer; however, some physical interaction is required of the musician, and as a consequence, the practice is not entirely freed from physical constraints. For instance, the musician can move his or her hands at only finite speeds. Because a musician can control live sound parameters in a live context, the keyboard connected to a loudspeaker is shown in Figure 5.2 fairly far to the top and to the right. However, the author personally does not consider the human-keyboard interaction to allow performances which are as expressive in their limit as the acoustic musical instrument—the musician is restricted by the physical parameters measured by the keyboard device. For instance, if the musician were to set the keyboard on fire, the result might be largely inaudible to the audience.²

Finally, a robot³ is shown playing a cello in the bottom leftmost corner of Figure

²We could similarly argue that an acoustic musical instrument’s vibration is described by an infinite number of modes of vibration. The dynamic behavior of these modes can be altered in an infinite number of diverse ways by the musician touching the instrument in an infinite set of positions. Consequently, we consider that the musician playing an acoustic musical instrument controls an infinite number of physical parameters, while the musician playing a digital interface controls only a finite number of measured parameters. Along these lines, we argue that acoustic musical instruments may in some sense allow for more expressive performances in the limit as the musician can theoretically leverage more playing techniques.

³This example image shows Pau de la Fusta, which is one of Carlos Corpa’s instrument-playing




Figure 5.3: Two extremes: a robot playing a musical instrument (left) and a human playing a musical instrument (right); and the spectrum in between (arrow)

5.2. The robot runs a program instructing it how to produce sound. The robot’s behavior can be made very precise because the program describes specifically what the robot should play. In addition, the live interaction is indeed very physical as the robot plays the acoustic instrument in a real setting.

5.4.2 Who Is In Control?

This thesis aims to make contributions in the applications of feedback control to musical instrument design. As a consequence, we consider the question, “**Who is in control for these musical interaction paradigms?**” Figure 5.2 suggests that we may largely ignore the vertical axis. Indeed, if we were to collapse the vertical axis, then some paradigms would overlap. The programming and the robotic musical instrument paradigms would overlap, while the digital interface and the traditional interaction paradigms would overlap, where the musician is in control. In other words, this question relates most directly to the precise control/live expressivity axis.

We continue the discussion now considering only whether a robot (computer) is robots from his fascinating Automatic Noise Ensemble.

in control or whether a musician is in control and refer to Figure 5.3. The right-hand side of Figure 5.3 depicts a musician playing an acoustic musical instrument in the traditional scenario, where the musician has intimate real-time control over the sound produced by the instrument. In contrast, the left-hand side of Figure 5.3 shows a robot playing an acoustic musical instrument. In the remaining chapters of this thesis, we study the middle ground between these two extreme practice paradigms. We design robots to assist the musician in making gestures. The musician should retain intimate control while gaining new advantages provided by robotics, such as speed and/or accuracy.

5.4.3 Philosophy

The first question to ask is, “At which end of the spectrum shown in Figure 5.3 do we find paradigms for useful musical practice involving human-robot interaction?”

Haptic Guidance

Previous experiments provide us with some insight into the left-hand side of the spectrum in Figure 5.3, where primarily the robot is in control. The robot essentially forces the musician to play a certain piece precisely. This scenario is referred to by many researchers as *haptic guidance*.

Graham Grindlay studied whether haptic guidance could be useful in teaching subjects pieces that a feedback controller forced them to play. Subjects’ recall of simple rhythms learned using audio-based training alone was compared against learning using audio-based training plus haptic guidance. Gridlay found that the addition of haptic guidance only mildly improved the subjects’ recall [72].

Outside the application of music performance, haptics researchers have had varying success in conducting related experiments involving haptic guidance. Feygin et al. showed that haptic guidance helped subjects learn to reproduce spatiotemporal trajectories [61]. Morris et al. showed that haptic feedback helped subjects learn to complete a surgery-similar task in which a sequence of forces was to be exerted along

a curve [112]. However, Gillespie et al. found that haptic guidance did not cause participants to learn a motor control skill faster [68]. Similarly, Yokokohji et al. found no significant benefit to haptic guidance in learning a very simple cube manipulation task [175]. Due to the mixed results in the field of haptic guidance and the difficulty in quickly conducting experiments that provide directly applicable results, we do not focus our research on haptic guidance.

Haptic Assistance

Instead we focus on another application of feedback control which we believe will yield more fruit. Specifically, we choose to operate closer to the right-hand side of the spectrum shown in Figure 5.3, where the musician plays the instrument and the feedback controller merely provides useful assistance. That is to say, in the haptic assistance performance paradigm, the musician has more control over the instrument than the robot. For example, Figure 5.4 shows a musician assisted by two SensAbleTM PHANTOM® Desktop haptic devices⁴. Each device exerts forces on one end of the bow, allowing the haptic devices to assist the musician in play. The haptic devices never force the musician to make a certain gesture; rather, they exert control only to influence the gestures that the musician makes.

Haptic Assistance Can Make Instruments Easier To Play

We would like to comment on the utility of haptic assistance: the robot can help the musician play the instrument. In other words, haptic assistance can make an instrument easier to play. Imagine a violin with haptic assistance, which helps the musician operate the bow correctly. Such an assistive violin would allow new players to skip the infuriating stage of initial practice where the bow makes screeching sounds. Other people living in the same house as the practicing musician would appreciate the assistive violin as well!

Haptic assistance is even more beneficial for novel musical instruments, which usually need to be learned quickly. These instruments typically have shorter lifespans

⁴PHANTOM, PHANTOM Desktop, SensAble, and SensAble Technologies, Inc. are trademarks or registered trademarks of SensAble Technologies, Inc.



Figure 5.4: A musician is assisted by haptic devices

and are played by few performers. Any reduction in the amount of practicing time for performers is beneficial. An additional side benefit of haptic assistance is that we can directly measure how well musicians perform subject to haptic assistance.

5.4.4 Human Motor Control

When feedback control is applied in a particular setting, the object to be controlled is usually modeled mathematically. A model makes it possible to estimate how the object will respond to feedback control. In many cases, the model is so precise that it can be used to derive optimal controller parameters [65]. In our setting, we control a musician's gestures. To model the musician, we can consult the models of the human motor system [139][142]. Unfortunately, these models are very complicated, so it would be difficult to make use of precise modeling for optimal control that would be practical in a real world music performance context.

While the musician's brain is too complicated to currently model well, we can at

Feedback mechanism	System delay time
Myotatic reflex	30ms-50ms
Long-loop reflex	50ms-80ms
Reaction time	120ms-180ms

Table 5.1: System delay times for the most important active feedback mechanisms of human motor system

least study the musician’s motor control system. This system has the ability to sense interactions with the environment, and the human motor control system can actuate the environment using the muscles. Various active feedback mechanisms exist that couple these sensing and actuation systems. The system delay times for the most important of these mechanisms are given in Table 5.1.

The fastest active feedback mechanism is the myotatic (stretch) reflex. The most familiar example of this reflex is the “knee-jerk reflex” triggered 30ms-50ms after the patellar tendon is struck with a small hammer [139]. This reflex mechanism is involuntary, so in order to assist the musician without triggering this type of reflex, we need to limit the amount by which we stretch the musician’s muscles.

The long-loop reflex, which is the second-fastest active feedback mechanism, is a reflex whose strength can apparently be modulated by commands from the central nervous system. This feature is believed to allow humans to temporarily resist strongly, for example while being tackled in American football, or to resist less, as when skiing [139]. The delay associated with the long-loop reflex ranges from 50ms to 80ms since the signals must travel along the spine to the thalamus and back [142].

The reaction time (RT) describes the time interval from the sudden presentation of an unanticipated force signal to the beginning of the human motor response, as directed by a signal returning from the brain [139]. RT is considered in more detail in the following section.

5.4.5 Closed-Loop Versus Open-Loop Control

Both closed-loop and open-loop control models provide distinct perspectives from which the human motor system can be studied. Humans use *closed-loop control* for

completing fine tasks such as threading a needle, where the control loop is closed around decision-making centers in the brain. The sequence of actions is as follows: the human attempts to move the end of the thread, waits for feedback from the senses about the new position of the thread, attempts to refine the position of the thread, waits for sensory feedback again, etc. The speed of closed-loop control of the human motor system is affected by the RT of the human motor system. Due to the relatively long motor system RT of roughly 120ms-180ms (see Table 5.1), the RT feedback mechanism is apparently not used for controlling brief tasks, such as for playing trills at up to 20Hz [120]. The reaction times for other modalities (auditory, visual, etc.) have similar durations, so they cannot be used to significantly speed up motor system responses either.

According to the *open-loop control* theory of the human motor system, once a human makes the decision to carry out a brief movement, a motor program is called up, which sequentially issues commands to the muscles. The way in which the commands are issued may depend on feedback, but the feedback in this loop is limited by the RT. For example, Muhammad Ali's left jab typically lasted only 40ms, and the essential part of a bat swing may take only 100ms [139].

There is evidence that many fast portions of tasks involved in playing a musical instrument, such as rapidly pressing a sequence of keys, are governed primarily by open-loop control. For instance, Schmidt discusses the

“example of a skilled pianist playing a piano with a broken key that could not be depressed. As the pianist played a string of notes, the attempts to press the broken key did not interrupt the series of actions at all. In fact, only after the entire sequence was completed did the individual notice and remark that the key was broken” [139].

Since it took the pianist longer to notice the broken key than the duration of a single note in the sequence, it seems likely that the pianist must have been playing the broken key using some form of open-loop control. Skilled typists [143] and telegraphers [38] are also believed to use open-loop control for completing motor tasks rapidly [120].

5.4.6 Haptic Assistance In General

If an assistive haptic interface ever takes an unexpected action too quickly, then a musician who uses some elements of open-loop control will not be able to respond appropriately, possibly making a mistake. Hence, we recommend,

Effective haptic assistance should never take an unexpected action so fast that it surprises the musician, which could prevent him or her from responding appropriately.

In practical contexts, we recommend more generally:

If haptic assistance is deterministic and relatively simple, then it is more likely to be useful.

5.4.7 Haptic Assistance For Playing Rhythm And Pitch

In this section, we argue that for our initial research into haptic assistance, we should investigate assisting musicians in playing rhythm and pitch. We formulate the argument by attempting to suggest what some of the most important features in music may be. To this end, we try to describe the gestures’ end result: the sound of the music corresponding to the gestures. Defining or even describing music in general is a difficult task—the Grove Music Online Dictionary’s entry entitled “music” is currently 22 pages long and contains many definitions as of this writing [116]. We make use of the definition originally taken from the Brockhaus-Wallring German Dictionary, which describes music as “die Kunst, Töne in ästhetisch befriedigender Form nacheinander und nebeneinander zu ordnen, rhythmisch zu gliedern, und zu einem geschlossenen Werk zusammenzufügen.” Translated into English, we define music as the following:

“The art of rhythmically organizing tones in succession and simultaneously in an aesthetically pleasing form and integrating them into a completed work.”

This definition is somewhat narrow, but it is congruent with many traditional definitions of music in that it describes music in terms of rhythm and tones.

The definition suggests that we should first consider assisting the musician in playing *rhythm*. In other words, we should develop feedback controllers that control the timing of events triggered by the musician. The definition also suggests that we should consider assisting the musician in playing tones. A musical tone can be described as having a distinct pitch, loudness, and timbre [59], but we do not wish to consider assisting in all of these different aspects of playing a tone. Hence, we consider the single tone characteristic that can often be the most distracting if played incorrectly in a piece: the *pitch*. In contrast, if a musician plays a tone or a group of tones with the incorrect loudness or timbre, the problematic result is generally less noticeable from a psychoacoustic perspective in many traditional settings—for one thing, the series of partials is usually less inharmonic [167][59].

In summary, we believe that some of the most fundamental forms of haptic assistance for musical performance revolve around assisting a musician in playing *rhythm* and *pitch*.

5.5 Overview

In this chapter we have explained our philosophy on haptic assistance and provided background information on related prior research on assistive haptic musical instruments. For suggestions on how to implement haptic musical instruments in practice, see the references by Sinclair and Cadoz as well as Appendices A and C.

Haptic assistance can

1. enable a musician to make new gestures that would otherwise be difficult or impossible, and to
2. aid a musician in playing more accurately.

We demonstrate (1) with respect to rhythm in Chapter 6 in the context of the Haptic Drumstick and Haptic Drum musical instruments. We demonstrate (2) with respect to rhythm in Chapter 6 and with respect to pitch in Chapter 7 in the context of

a formal subject test with musicians. The musicians perform according to different force conditions on a haptic Theremin-based instrument.

Chapter 6

Assisting Musicians in Making New Gestures That Would Otherwise Be Difficult or Impossible

Haptic assistance can aid musicians in making new gestures that would otherwise be difficult or impossible. We demonstrate this application of haptics with respect to controlling rhythm so that musicians can leverage new physical interactions that are immediately intuitive. To this end, we introduce the following two interfaces:

- the Haptic Drumstick, which is a haptically augmented drumstick that strikes a *virtual* surface, and
- the Haptic Drum, which is a haptically augmented but completely *physical* drum pad.

Even though the Haptic Drumstick and Haptic Drum are constructed using different haptic devices, they result in similar dynamic behavior. They assist a musician in single-handedly playing drum rolls. The Haptic Drumstick and the Haptic Drum can enable the musician to play drum rolls which are especially fast, have particular rhythmic characters, and/or are arbitrarily complex.

In traditional drumming, damping removes energy from the drum stick with each bounce. An effective control algorithm is to apply a pulse to the drumstick every time that the drumstick strikes the drum membrane. When the musician presses down harder with the drumstick, the frequency of the drum roll increases. We analyze the stability of the control system using a Poincaré map.

6.1 Simple Model of Drumstick/Membrane Interaction

First we introduce a simplified model describing the traditional drum stick/membrane interaction. We develop the simplest model that can describe how drummers are able to play drum rolls and which allows implementation of the same interaction on a haptic device.

Consider the double stroke roll from traditional drumming. The drummer throws a stick at the membrane, allows it to bounce twice, retracts it, and then repeats the action with the other stick. The manner in which the drum membrane throws the stick back toward the drummer helps facilitate both the bouncing and retracting actions.

Because we are concerning ourselves primarily with relatively fast drum rolls, we assume that the angle of rotation of the drumstick θ does not change much (see Figure 6.1). This simplification allows us to linearize the rotation of the drumstick tip, and so we model the vertical motion of the drumstick tip as a bouncing point mass m .

6.1.1 Above The Drum Membrane

The hand grasping the stick acts as both a rotational spring and a rotational damper at the butt of the stick (not shown). For simplicity, we linearize these elements and commute them to the drumstick tip, representing them by K_{hand} and R_{out} (see Figure 6.1). By changing the grasp of the stick, the drummer can change K_{hand} and R_{out} . This is known as *passive impedance modulation* and allows drummers to play drum rolls at rates up to 30Hz, even though the human neuromuscular system has a reaction




Figure 6.1: Drumstick dynamics for drumstick above membrane

time of over 100ms [77]. By considering that the drummer also may exert a slowly-varying, but for our purposes, approximately constant force on drumstick, we realize that we must also include an approximately constant bias force term F_{bias} . We lump the effects of gravity into F_{bias} . We write the equation of motion in the absence of collisions with the drum membrane (i.e. for vertical displacements $z > 0$):


$$m\ddot{z} + R_{out}\dot{z} + K_{hand}z = F_{bias} \quad (6.1)$$

6.1.2 Simple Collision Model

The physics-based approach to modeling stick/membrane collisions involves the coefficient of restitution (COR) β . For a collision beginning at any time t_{in} and ending at any time t_{out} , if $\dot{z}(t_{in}) = v_0$, then $\dot{z}(t_{out}) = -\beta v_0$. On the other hand, in haptics, the standard method for modeling a collision involves a spring-like penalty force implemented by a spring with constant $K_{coll} \gg K_{hand}$ [137]. There must also be some damping factor R_{in} due to losses absorbed by the hand and the collision, so we arrive at the dynamics for $z < 0$ (see Figure 6.2):

$$m\ddot{z} + R_{in}\dot{z} + K_{hand}z + K_{coll}z = F_{bias} \quad (6.2)$$

The term F_{bias} causes the COR β to become weakly dependent on the velocity $\dot{z}(t_{in})$ at the beginning of a collision. However, since the collision is quick, the terms involving R_{in} and K_{coll} dominate. Neglecting the other terms, we may solve the differential

Figure 6.2: Drumstick dynamics for $z < 0$

equation analytically to arrive at

$$\beta \approx \exp \frac{-R_{in}\pi}{\sqrt{4mK_{coll} - R_{in}^2}}. \quad (6.3)$$

6.1.3 Upward Soft Collisions

Noting the similarity between (6.1) and (6.2), we suspect that we may also determine a COR α for the “soft collision” against the hand in the upward direction. However, since the spring constant K_{hand} is relatively small, the collision takes longer, and so we can no longer neglect F_{bias} . Consequently, $\alpha(\dot{z}(t_{out}))$ is strongly dependent on the velocity $\dot{z}(t_{out})$ at the beginning of each “soft collision,” so the expression for α is more complicated than the expression for β . Nevertheless, in simulations of (6.1) for realistic system parameters, we have verified that α remains roughly constant for bounces of similar amplitude.

6.2 Haptic Drumstick

6.2.1 Overview

The model from the previous section is used to implement the Haptic Drumstick on a high-resolution haptic device. The musician grips a physical stick, but the drum is entirely virtual. We create a new musical instrument by altering the haptic drumstick dynamics to facilitate similar yet new physical interactions. We focus on drum rolls. In particular, we alter the haptic drumstick dynamics to assist performers in playing single-handed drum rolls [27]. Finally, we analyze the stability of the altered system dynamics using a Poincaré map.

6.2.2 Introduction

To construct the haptic drumstick we desire an appropriate haptic display for implementing the equations of motion. Haptic drumsticks have previously been implemented using single DOF haptic displays [13].¹ We use the three DOF Model T PHANTOM haptic device² because it has relatively strong motors—for moderately strong strokes, the motors are able to render stiffnesses comparable to a typical drumhead. The performer holds the drumstick-like stylus in one hand as shown in Figure 6.3.

The sound synthesis engine consists of a digital waveguide mesh modeling waves propagating in a drum membrane [149]. The force of the drumstick on the spring K_{coll} is fed into the nearest node in the mesh. This way, striking the modeled drum at different positions results in different sounds. Figure 6.4 shows how the elements of the system are interconnected. The PHANTOM sends the vertical displacement z to a block that calculates F_m , the force to be exerted on the PHANTOM. F_m is also passed to the digital waveguide mesh in order to facilitate sound synthesis. A video demonstration is available online.³

¹<http://web.media.mit.edu/~grindlay/FielDrum.html>

²From SensAble Technologies, see <http://www.sensable.com>.

³<http://ccrma.stanford.edu/~eberdahl/Projects/HapticDrumstick>




Figure 6.3: Graphical display (above) and PHANTOM robotic arm (lower right)




Figure 6.4: Block diagram describing the operation of the haptic drumstick

6.2.3 Altering The Physical Interaction

There are many ways in which the physical interaction could be altered. We consider alterations involving noise and deterministic chaos to be less desirable because humans do not typically encounter these physical interactions in nature. In contrast, stable limit cycles, which are self-sustaining attractive oscillations, describe the behavior of biological oscillators such as the heart. Limit cycles also manifest themselves in bowed strings, vibrating reeds, and drum rolls [62].

Most drummers use two hands to play drum rolls. This limits the types of patterns that they can play. However, by inducing limit cycle behavior, we make it easy for a drummer to single-handedly play a drum roll. In informal tests, we determined that implementing system delay, a hysteretic spring K_{coll} , or negative damping R_{out} make the haptic drumstick more likely to self-oscillate, but not in a particularly intuitive way. Indeed, such mechanical elements do not occur readily in nature. A more intuitive solution involves forcing the drumstick in the positive z -direction by the pulse $h(t)$ every time the stick enters the simulated membrane, where

$$h(t) = \frac{m\Delta v_{pls}}{\tau} e^{-t/\tau} \quad (6.4)$$

and $\tau \approx 2$ ms. These force pulses are superimposed with the forces described by (6.1) and (6.2). Assuming a quick enough pulse, the pulse $h(t)$ causes the velocity of the drumstick tip to change by Δv_{pls} . Playing a drum roll according to this physical interaction feels somewhat like holding a stick against the spokes of a rotating bicycle wheel. This control algorithm is an example of event-based control (see Section 4.3). A similar algorithm has been employed for improving the stability and perceived rigidity of virtual surfaces rendered using haptic interfaces [50][101].

Directly Altering The COR

Since we can effectively change the velocity of the drumstick after the collision, by choosing $\Delta v_{pls} = -\gamma \dot{z}(t_{in})/\beta$ for some $\gamma > 0$, we can obtain a new COR $\hat{\beta} = \beta + \gamma$:

$$\dot{z}(t_{out}) = -\beta(\dot{z}(t_{in}) - \frac{\gamma}{\beta}\dot{z}(t_{in})) = -(\beta + \gamma)\dot{z}(t_{in}) = -\hat{\beta}\dot{z}(t_{in}) \quad (6.5)$$

With $\hat{\beta} > 1$, we can counter the damping due to the hand.

Applying Pulses With Constant Magnitude

Choosing Δv_{pls} to be constant presents a superior alternative because the total energy in the system becomes limited. This safety mechanism prevents drummers from inadvertently damaging the PHANTOM robotic arm because the self-oscillations tend to be small. Consider that for vibrations at small amplitudes, the effective $\hat{\beta}$ is large, and that for vibrations at large amplitudes, the effective $\hat{\beta}$ is approximately equal to β .

Hence, the force on the drumstick F_m is given by

$$F_m = -K_{coll}z \cdot (z < 0) + l(t), \quad (6.6)$$

where $l(t)$ describes the pulses with constant magnitude. For instance, if there are N collisions total, then there are N pulses total and

$$l(t) = \sum_{i=1}^N h(t - \tau_i), \quad (6.7)$$

where τ_i represents the i th time that z becomes negative. In the next section, we show that choosing the pulse magnitude Δv_{pls} constant leads to stable limit cycle behavior.

6.2.4 Stability Analysis

Analysis

We analyze the altered dynamics with a Poincaré map, which allows the stability of a closed orbit in a continuous-time system to be determined from the stability of a related discrete-time system [138]. Let $\mathbf{z} = (z, \dot{z}) \in \mathbb{R}^2$ describe the system state, and let the *system flow* $\phi_t(\mathbf{z}_0)$ describe the current state given an initial state \mathbf{z}_0 at t seconds in the past. Call d the periodic orbit of the system in the phase plane of \mathbf{z} , and consider the semi-infinite line E , which is crossed once per cycle by system




Figure 6.5: Example system trajectories in the phase plane

trajectories (see Figure 6.5).

$$E = \{(z, \dot{z}) | z = e, \dot{z} > 0\} \quad (6.8)$$

For simplicity, we take $e > 0$ to be arbitrarily small. We assume that $F_{bias} < 0$ so that after leaving the membrane, the drumstick will eventually strike it again. To simplify the analysis, we restrict U to be a small enough neighborhood of p on E so that α is approximately constant (see section 6.1.3). We further assume that the pulse $h(t)$ is short enough that it always ends before the trajectory intersects E . If we then define the map $P : U \rightarrow E$ so that for $q \in U$

$$P(q) = \phi_{\tau(q)}(q) \quad (6.9)$$

where $\tau(q)$ is the time for the orbit $\phi_t(q)$ based at q to first return to E , $P(\cdot)$ is a Poincaré map. Consequently, we may analyze the stability of the closed orbit d by analyzing the stability of the discrete-time system $P(v_i)$, where v_i is the velocity of the drumstick tip at the end of the i th collision with the drum membrane.

$$v_{i+1} = P(v_i) = \alpha\beta v_i + \beta\Delta v_{pls} \quad (6.10)$$

Directly Altering The COR

For the alteration where $\Delta v_{pls} = -\gamma \dot{z}(t_{in})/\beta = \gamma \alpha v_i/\beta$,

$$v_{i+1} = P(v_i) = \alpha(\beta + \gamma)v_i. \quad (6.11)$$

If the alteration of the dynamics is configured with γ such that $\alpha(\beta + \gamma) = 1$, then (6.10) describes a marginally-stable system, so the drumstick will oscillate at constant amplitude. However, this behavior is not stable—any parameter deviation will lead to a decaying oscillation or a growing oscillation. Since $P(\cdot)$ is a Poincaré map, the closed orbit d is *not a stable limit cycle*. Nevertheless, a drummer may stabilize this system by adjusting α in real time using passive impedance modulation.

Applying Pulses With Constant Magnitude

When Δv_{pls} is held constant, (6.10) describes a discrete-time linear system driven by a constant input. With the exception of the pulse $h(t)$, collisions with the membrane and upward soft collisions against the hand are dissipative. This means that $\alpha < 1$ and $\beta < 1$. Then $\alpha\beta < 1$, so (6.10) describes a stable discrete-time system. Since $P(\cdot)$ is a Poincaré map, the closed orbit d is a *stable limit cycle*. We can also use (6.10) to easily calculate other properties of the system. Since Δv_{pls} is constant, v_i will approach the steady-state $v_{lc} = \frac{\beta \Delta v_{pls}}{1 - \alpha\beta}$.

6.2.5 Results

In informal tests, we found that the altered dynamics made playing single-handed drum rolls easy for both 1) directly altering the COR and 2) applying pulses with constant magnitude. The main advantage of (2) is that drum roll limit cycles are guaranteed to be stable. However, both types of altered dynamics allow drummers to increase the drum roll rate by increasing K_{hand} or decreasing F_{bias} as in traditional drum roll playing [77]. This means that the new physical interaction is similar to the physical interaction when playing traditional drums. Consequently, the Haptic Drumstick should be easy to play because musicians can leverage prior experience in

playing drums. Informal testing suggests also that musicians do indeed find it easy to learn to play the Haptic Drumstick.

6.2.6 Limitations

In practical performance contexts, the Haptic Drumstick suffers from some limitations imposed by the physical limits of the PHANTOM haptic device. Many drummers wish that the Haptic Drumstick could render the large forces that they expect. Instead, if a drummer strikes the virtual drum hard enough, the motor currents saturate, allowing the Haptic Drumstick to “push through” the virtual drum surface rather than pushing back outward. It is also possible that the haptic device motors could overheat due to the large currents passing through them during attempts to repeatedly render strong opposing forces. While Bennett addresses some of these issues by introducing an MR-fluid brake as the actuator for a haptic interface to a virtual drum, his device could not implement the type of haptic assistance that we describe here as it cannot add energy to the system [13].

In addition, drummers expect to be able to move a drumstick within a volume extending as far as they can reach. However, the Haptic Drumstick is limited to the relatively small workspace allowed by the haptic device. Furthermore, many drummers prefer using certain drumsticks because of the way in which they are weighted. To solve these problems, we introduce a new device in the next section, where a drum pad is haptically augmented.

6.3 Haptic Drum

Making the drumming surface physical rather than purely virtual mitigates the aforementioned limitations. In this section, we describe the design of a new interface called the haptic drum. We present some measurements of its dynamic behavior.




Figure 6.6: Haptic Drum v1

6.3.1 Prototype v1 Design

The Haptic Drum is designed around a woofer actuator because to some extent, a woofer physically resembles a drum membrane.⁴ Most importantly, a woofer behaves like a spring at low frequencies, so no control energy is required to emulate the stiffness of a drum membrane. A woofer can be relatively small and operates in only one degree of freedom; however, this is sufficient for implementing many of the interesting dynamical behaviors involved in haptically assisted drumming.

The basic design is shown in Figure 6.6 (left). A woofer is attached to a piece of white wood, which hides the stand-alone controller circuit from view. Three black knobs can be used to change the controller parameters. In order to protect the woofer, the end of a coffee can is glued to the woofer cone and protrudes up through a large hole cut in the white piece of wood. With a drumstick, the musician strikes the end of the coffee can, which serves as a drum pad as shown in Figure 6.6. Five piezoelectric disks are glued to the underside of the drum pad in order to detect the approximate position of drumstick impacts (see the orientation in of the disks in Figure 6.6, right). The controller can optionally be connected to an external computer for synthesizing sound in response to drumstick impacts.

Because the Haptic Drum operates based on the principle of real mechanical vibrations, it can be activated using implements besides drumsticks. Figure 6.7 shows

⁴The reading committee agrees that the work for the Haptic Drum and Haptic Drumstick was not part of Edgar Berdahl's University responsibilities in the context of the SU-18 agreement.



Figure 6.7: Haptic Drum prototype v1 with additional percussion instruments as used for performing “Percussion Edstravaganza”

the author playing the Haptic Drum with an egg shaker and a drumstick simultaneously. When the egg is dropped upon the drum pad, the egg bounces around due to the feedback control. The drumstick can be employed to modify the manner in which the egg bounces. Additional percussion elements, such as a splash cymbal, a tambourine, and another shaker can also be used along with the Haptic Drum in performance (see Figure 6.7).

6.3.2 Dynamic Behavior

Control Algorithm

The most basic Haptic Drum v1 control algorithm applies pulses with constant magnitude as explained in Section 6.2.4. It is adapted to the Haptic Drum hardware as follows:

1. Wait until a drumstick impact is detected.
2. If a USB cable is connected to the device, send an output signal describing the impact to allow for sound synthesis by an external device.
3. Send a pulse Δv_{pls} with fixed magnitude to the loudspeaker to cause the drumstick to bounce upward.
4. Wait for M milliseconds while the cross-talk from the actuated pulse Δv_{pls} to the sensors decays.
5. Go back to step 1.

Basic Measurements

The dynamic behavior once a drumstick strikes the drum pad is to first order the same as the behavior described in Section 6.2.4. We present measurements of a drumstick's motion during interaction with the Haptic Drum v1. The vertical position of the drumstick is measured over time with the Fairchild QRB1114, an optical reflective object sensor. The zero position is calibrated to be the position of the drum pad in the absence of play. However, since the woofer has its own dynamics, collisions between the drum pad and the drumstick do not always occur precisely at the zero position.

Due to the biasing of the sensor, we did always know the magnitude of the position variations precisely. For measurements involving large position deviations, we approximately calibrated the vertical units visually. For measurements involving small position deviations, the position units are not known, so they are plotted without




Figure 6.8: Vertical position of the drumstick while a musician plays a regular drum roll at 70Hz

units in that case. For the uncalibrated measurements, the position likely varies over the order of a few millimeters.

Fast Drum Rolls

Without haptic assistance, musicians can play drum rolls at rates of up to 30Hz [77]. The Haptic Drum can enable a musician to play even faster drum rolls. For instance, the position trajectory shown in Figure 6.8 shows the position of the drum stick as the musician plays a very regular drum roll at about 70Hz on the Haptic Drum v1. An impact between the drumstick and the Haptic Drum occurs every time that the position stops decreasing and starts increasing (see Figure 6.8).

Complex Drum Rolls

The Haptic Drum can assist the musician in playing arbitrarily complex drum rolls that would not be possible without haptic assistance. Figure 6.9 illustrates one example where a second drumstick is also bouncing on the Haptic Drum. Its impacts (not shown) interfere with the impacts of the first drumstick whose position is shown in Figure 6.9 since only one discrete impact can be detected during any window of M milliseconds, as explained in Section 6.3.2. The pressure applied via the second drumstick changes the the average height of the drum pad (see the time-varying level of the drumstick impacts in Figure 6.9).




Figure 6.9: Vertical position of the drumstick shown for complex behavior

Example Phase Plane Trajectories

To verify that the Haptic Drum implements similar dynamics as the Haptic Drumstick, we present a measurement of a drumstick in the phase plane. Figure 6.10 (top) shows the position of a single drumstick playing a roll at about 5Hz on the Haptic Drum v1. An estimation of the velocity, which is shown in Figure 6.10 (bottom), reveals that as the drumstick leaves the Haptic Drum pad, its velocity is positive and large, then the velocity gradually becomes negative, causing the drumstick to eventually strike the Haptic Drum pad again.

The velocity is plotted as a function of position in Figure 6.11 for two impacts. As time moves forward, the (position, velocity) point moves clockwise around the loop, which represents a stable limit cycle. The same data is used as in Figure 6.10. Each time around the loop, the trajectory changes slightly as the stiffness of the hand and the force with which the hand presses the drumstick downward vary with time.

As expected, the phase plane trajectory is similar to the simplified Haptic Drumstick trajectory shown in Figure 6.5.

1. In the case of the Haptic Drum, the pulse delivered by the Haptic Drum dominates the stiffness presented by the mechanical properties of the drum pad constructed out of a coffee can. The pulse is delivered during the brief instant when the position is approximately zero (see Figure 6.11).
2. The damping in the hand causes the speed with which the drumstick departs




Figure 6.10: Drumstick position (top) and velocity (bottom) for a low-rate drum roll

the drum pad to be greater than the speed with which the drumstick arrives at the drum pad (see Figure 6.10, bottom).

The Haptic Drum can also allow the musician to single-handedly play a roll that resembles the double stroke roll. In this case, the control algorithm alternatively responds to each impact with a stronger and a weaker actuation pulse Δv_{pls} . When the pulse is stronger, the delay is longer before the drumstick impacts against the drum pad again. Four example impacts are shown in Figure 6.12. The corresponding (position, velocity) phase plane trajectory is shown in Figure 6.13.

6.3.3 Prototype v2 Design

A second prototype is appropriate for more traditional music making applications. It employs a larger woofer and a more robust drum pad so that drummers can hit the Haptic Drum v2 as hard as they hit normal drums without damaging it. The second prototype is constructed within a small snare shell, as shown in Figure 6.14, to emphasize the fact that the haptic drum can also be used in standard drum sets alongside standard drums. An extra small drum synthesizer from Roland is built into the prototype to allow the timbre of the drum to be changed without any external




Figure 6.11: Example phase plane trajectory for two impacts corresponding to Figure 6.10




Figure 6.12: Drumstick position (top) and velocity (bottom) for a single-handed haptically assisted double stroke roll




Figure 6.13: Example phase plane trajectory for four impacts corresponding to the data shown in Figure 6.12

hardware (see Figure 6.14, top). Figure 6.15 depicts a musician playing the Haptic Drum v2 when set up with a normal drum set. A promotional flyer is shown in Figure 6.16.

6.4 Conclusions

Stable limit cycles manifest themselves not only in drum rolls but also in bowed strings, vibrating reeds [62], self-sustaining electric guitar strings, the heart, and other biological oscillators [138]. As a consequence, limit cycle-based physical interactions are not only physically intuitive but also attractive to musicians. Even though the Haptic Drumstick and Haptic Drum are based on different haptic devices, they both tend to induce limit cycle-type behavior. These devices can enable the musician to play single-handed or two-handed drum rolls which are especially fast, rhythmically characteristic, and/or arbitrarily complex.

Both the haptic drumstick and the haptic drum are non-passive: they add energy to the drumstick in order to allow it to roll for an indefinite period of time. This is one of the reasons why we include the guideline: “many useful haptic assistance systems are not passive” among our haptic assistance guidelines in Section 8.6.



Figure 6.14: Haptic Drum prototype v2 with built-in synthesizer



Figure 6.15: Haptic Drum prototype v2 used with standard drum set

HAPTIC DRUM

Berdahl Innovations

The word *haptic* comes from Greek and pertains to the sense of touch. The haptic drum harnesses the power of force-feedback to assist drummers in playing parts that would otherwise be difficult or impossible. This patent-pending device consists of a drum pad, a DSP, an amplifier, and a woofer. Whenever a drumstick impacts the drum pad, the woofer gives a small push in the upward direction, adding energy to the bouncing drumstick.




The haptic drum

- makes it easy to play a one-handed drum roll, freeing up the other hand
- can sustain drum rolls for arbitrarily long periods of time
- allows drum rolls to be played at very fast rates (e.g. up to 60 impacts/second)
- makes it possible to play exotic new drumming rudiments
- is compatible with the standard drum set
- fits inside of a 10" snare drum shell

Quote:

"Now electric guitarists won't be the only ones having all the fun with their electronic effects!"

Video demonstration at: <http://ccrma.stanford.edu/~eberdahl/Projects/HapticDrum>



Come see what all the fuss is about!
Contact Edgar Berdahl for a demo.

Edgar J. Berdahl

Figure 6.16: Promotional flyer for the Haptic Drum v2

The Haptic Drumstick and Haptic Drum demonstrate that haptic assistance can aid musicians in making new drum roll and percussion rudiments gestures that would otherwise be difficult or impossible. Other forms of haptic assistance will be useful in aiding musicians in making other types of difficult gestures.

Chapter 7

Assisting Musicians in Making Gestures More Accurately

Haptic technology, providing force cues and creating a programmable physical instrument interface, can assist musicians in making gestures more accurately. We begin by demonstrating that the Haptic Drum can assist a musician in playing a drum roll especially accurately. The remainder of the chapter is devoted to a subject test showing that haptic assistance can be used to help musicians select pitches more accurately on a Theremin-like interface.

7.1 More Accurate Rhythmic Elements

The haptic drum introduced in Chapter 6 can assist a musician in playing drum rudiments accurately. For instance, if the goal is to play a drum roll with a regular period, the Haptic Drum can enable a musician to do so—the musician need merely attempt to exert a constant downward force on the haptic drum with the drumstick, and the Haptic Drum v1 does the rest. Figure 7.1 illustrates an example involving the control algorithm where pulses are exerted with constant magnitude (see Section 6.2.4). The drumstick follows the same approximate trajectory repeatedly: the position is shown above and the velocity of the drumstick is shown below in Figure 7.1.




Figure 7.1: Drumstick position (above) and velocity (below) for a regular drum roll

7.2 More Accurate Pitches

We consider the specific problem of assisting a musician in selecting pitches from a continuous range. We build on a prior study by O’Modhrain of the accuracy of pitches selected by musicians on a Theremin-like haptic interface [120]. To improve the assistance, we augment the interface with programmed detents so that the musician is gently guided toward locations of equal tempered pitches. Nevertheless, the musician can still perform arbitrary pitch inflections such as glissandi, falls, and scoops. We investigate various forms of haptic detents, including fixed detent levels and force-sensitive detent levels.

We carry out a formal subject test with 16 participants to compare the efficacy of four different forms of haptic feedback. One control condition involves no haptic feedback, one condition involves simple spring force feedback, and two conditions involve haptic detents. Statistically significant results from the formal subject test confirm improved accuracy brought about by detents.

7.3 Pitch Selection Problem Overview

A musical tone can be described as having a distinct pitch, loudness, and timbre [59]. We study how haptic devices can assist musicians in selecting pitches accurately. Consider the different ways in which pitches are selected while playing the piano versus playing the cello. The piano has 88 strings of fixed length, so a pianist using standard playing techniques can play only a discrete set of pitches. If the strings are tuned consistently, then no matter which keys the pianist presses, the performance will be approximately in tune. In contrast, the length of a cello string can be adjusted directly by the musician's hand, meaning that a cellist can play pitches over a continuous range. This freedom can be troublesome for beginning cellists who might for example wish to initially limit themselves to playing a discrete set of pitches corresponding to a scale. In fact, this freedom can allow also vocalists, trombonists, violinists, viola players, string bass players, etc. to unintentionally play out of tune.

In a broad sense, we seek to combine the way that pitches are selected on a piano and on a cello in order to make it easy to play pitches accurately while still allowing arbitrary continuous pitch inflections to be created. More specifically, we consider using haptic feedback to improve the accuracy with which a musician can select a pitch over a continuous range. We design the haptic assistance to gently guide the musician to desirable pitches. Of course, musicians may disagree over which pitches are *desirable*. In fact, the pitch for some sounds can be ambiguous, so we need to be more specific.

For our work, we use tones consisting of sinusoids whose frequencies are integer multiples of a *fundamental frequency*. Hence, the fundamental frequency of the tones is mathematically well defined, and for the tones we use, the perceptual quantity, the pitch, is nearly the same as the fundamental frequency [59]. From this point onward, we will use only the term fundamental frequency when discussing the particulars of a given musical instrument or type of haptic assistance, and we will use pitch when discussing the general problem or the musician's perception of fundamental frequency.

We make another approximation by stating that we will assist the musician in playing the fundamental frequencies from the chromatic scale defined by the equal

temperament tuning system [59]. Most common Western tuning systems are relatively similar to equal temperament, and we will choose performance tasks whose fundamental frequency errors will be larger than the relatively small tuning ambiguities described here. We consider this, albeit limiting, discrete set of fundamental frequencies to be *desirable* for the purposes of studying haptic assistance in a scientific context.

7.4 Prior Work

The Theremin electronic instrument, which was patented in the United States in 1928, produces a harmonic tone as output [156]. The fundamental frequency of the tone is controlled by the position of one hand in free space, while the amplitude of the tone is controlled by the position of the other hand. The Theremin instrument provides no haptic feedback, and it is often informally considered to be difficult to play as it is non-trivial for the musician to orient his or her hands in free space.

O’Modhrain studied the accuracy with which musicians select fundamental frequencies with a Theremin-like interface implemented using a haptic device. She compared controlling fundamental frequency given the complete absence of haptic feedback versus several kinesthetic haptic feedback conditions such as a spring force, a viscous damping force, and a constant force. In Chapter 4 of her PhD thesis, she draws the conclusion that the “existence of force feedback in a computer-based instrument marginally improves performance of a simple musical [fundamental frequency selection] task” (p. 49) [120].

Besides Theremin-type interfaces, glove-based and other continuous interfaces may lack significant haptic feedback. We generalize O’Modhrain’s conclusion by hypothesizing that if a musical instrument does not provide any haptic feedback at all, it will probably be more difficult to play accurately. We term this hypothesis the “*Theremin Hypothesis*.” As a consequence, we recommend that musical instrument designers incorporate haptic feedback into their instrument designs.

We hypothesize further that specific kinds of active force feedback assisting the musician in selecting desirable fundamental frequencies may be even more effective.

This hypothesis is also suggested by Moss' and Cunitz's work in which a specific kind of haptic feedback pushes the musician's finger toward the notes of the chromatic scale; however, Moss and Cunitz do not consider any other types of haptic feedback [113].

7.5 Simple Theremin-Like Instrument

We have strived to make the example instrument in our laboratory as easy to play as possible. Although the PHANTOM Model T can move in three dimensions, we restrict motion of the thimble to the y -axis and simultaneously measure pressure applied normal to this axis in the vertical z -direction (see Figure 7.2). Rather than simply waving his or her hand around in space, the musician moves the thimble with his or her fingertip. While doing so, the musician can rest his or her wrist on the white piece of wood placed horizontally across the table as shown in Figure 7.2. For simplicity, we allow the musician to adjust only the fundamental frequency of the sound. Sound is synthesized using additive synthesis with sinusoidal components at the fundamental frequency, twice the fundamental frequency, and three times the fundamental frequency [134].

Due to the finite workspace of the haptic device, the horizontal position y (see Figure 7.2) of the musician's finger can vary over about 20cm and is mapped to the logarithm of the fundamental frequency of the musical instrument. This mapping allows the distance between each pair of adjacent notes in the chromatic scale to be about 0.75cm. Higher fundamental frequencies are further to the right, as with the pianoforte. However, in comparison with the pianoforte, each octave on the Theremin-like instrument spanned roughly half the distance due to the workspace constraints of the PHANTOM Model T.

We have experimented informally with other mappings, such as an affine mapping between position and fundamental frequency as well as the nonlinear mapping between the length of a vibrating string and its fundamental frequency. We concluded that the above-described logarithmic mapping, as suggested by O'Modhain




Figure 7.2: Musician's finger inserted into the PHANTOM thimble

in personal communication, is best not only because of the human's logarithmic perception of pitch, but also because of the equal spacing of the chromatic notes along the instrument interface.

7.6 Haptic Assistance For Fundamental Frequency Selection

The haptic assistance operates sideways from the perspective of the musician. In particular, we consider controllers that exert force in the lateral y -axis as a function of the current and past y -positions of the thimble (see Figure 7.2). Many forms of haptic feedback are imaginable, so we limit these forms here by considering the human-computer interaction (HCI) literature.

7.6.1 Basic Detent

The basic detent is simple, deterministic, and can help the musician orient himself or herself. A detent can be created even using 1DOF haptic interfaces. Figure 7.3 illustrates how to implement a simple piecewise linear detent. Near the center of the detent, the force in the lateral y -axis behaves like that of a spring, while the force




Figure 7.3: Force profile $F(y)$ and terrain height profile $h(y)$

goes to zero when the position y moves further from the detent's center [166]. The force profile is consistent with the lateral forces one would experience in the terrain height map $h(y)$ shown at the bottom of in Figure 7.3.

Researchers from the human computer interaction (HCI) literature have experimented with piecewise linear as well as piecewise nonlinear detents [113]. They report that altering the shape of the detent nonlinearly manifests itself psychophysically as apparently only an intensity difference [173][125]. We have also carried out some informal tests and also believe that the details of the detent shape are not of primary importance, so we believe the piecewise linear detent design to be sufficient for our study. For more information on how we tuned the parameters describing the exact shape of our basic detent, see Section 7.7.3.

7.6.2 Force Feedback Conditions

We now describe the specific haptic force feedback conditions that we incorporated into our subject test.

Multiple Detents (DET)

We extended the basic detent described in Section 7.6.1 to assist the musician in playing notes from a diatonic scale consisting of whole and half steps, which is based at the origin of the haptic device. Figure 7.4 shows the piecewise linear force profile for the first five notes of the scale. A spring force field was centered around each note, making the y -position a locally stable equilibrium point, as denoted by each dashed blue circle [136]. In between each pair of notes, the forces were tapered toward one




Figure 7.4: Force profile for the DET condition

another to avoid creating any distracting discontinuities. Since there was a half-step instead of a whole-step between “mi” and “fa”, the force profile was warped to retain the same form (see Figure 7.4).¹ Note that we made the detent force profiles as wide as possible to ensure that we were providing haptic assistance for the largest possible set of positions y .

Velocity-Sensitive Detents (VEL)

While the HCI literature was in agreement that a single detent centered around a single target, such as a menu item, improves performance, it was harder for interface designers to improve performance when faced with multiple possible targets, each having its own detent, especially when the user needed to traverse through distracting detents in order to reach a specific target. Researchers reported apparent improvements in performance by making the magnitude of the detents dependent on the velocity [118][119][117][124]. However, Hwang et al. found that no velocity dependence was necessary for motion impaired users to benefit from having multiple targets [89].

We originally hypothesized that by considering the notes to be the analog of menu targets, we might take advantage of this feature to improve upon the DET condition.

¹One might ask whether the force magnitude of the detents ought to be doubled for half steps, such as for the distance between “mi” and “fa”, in comparison with whole steps. This change would cause the amount of work required to move the thimble from one zero-crossing of $F_{DET}(y)$ to another to be constant. However, in informal tests, we found having the force magnitude double for some detents to be much more distracting. In fact, we found the force profile shown in Figure 7.4 to be much less distracting, so it is the profile that we suggest here.




Figure 7.5: Strength of haptic assistance for the VEL condition

We defined the feedback force condition $F_{VEL}(y, \dot{y})$ as follows:

$$F_{VEL}(y, \dot{y}) = F_{DET}(y) \cdot g(|\dot{y}|). \quad (7.1)$$

$g(\cdot)$ was designed so that little to no haptic assistance was provided for relatively large speeds $|\dot{y}|$, while significant haptic assistance was provided at slow speeds (see Figure 7.5).

After some informal experimentation with this force condition, we became of the opinion that some form of velocity-sensitive detents would likely be useful. Nevertheless, the usefulness of the force condition seemed to be sensitive to the shape of $g(\cdot)$. In addition, we were not entirely convinced that the benefit of VEL would be worth the effort of optimizing all the parameters describing it. For example, $g(\cdot)$ could have been described by a nonlinear function rather than by a piecewise linear function (see Figure 7.5). In the interest of conducting a statistically significant formal subject test that could be compared to O’Modhrain’s work [120], we chose to include the next force condition instead of VEL in the subject test.

Force-Sensitive Detents (FRC)

Past experiments by other researchers indicated that the presence of multiple detents could be problematic because users could “get stuck” in interfering detents [118][119]. We proposed a new solution, which we believe to be novel. We believed that the musician should be able to regulate the degree of haptic assistance in real time. In other words, rather than trying to infer from \dot{y} whether the musician required




Figure 7.6: Strength of haptic detents for the FRC condition

assistance, we desired that the musician be capable of specifying in real time how much assistance he or she was receiving. The simplest approach was to allow the musician to control the degree of haptic assistance with the downward force he or she applied to the thimble. We configured the FRC force condition so that the musician only received assistance when he or she requested it by pressing downward with force $p > 0$. If the musician instead pulled the thimble even slightly upward (i.e. $p < 0$), then no haptic assistance was provided at all, allowing the musician to breeze over any otherwise distracting detents. We defined $F_{FRC}(y, p)$ as follows:

$$F_{FRC}(y, p) = F_{DET}(y) \cdot S(p), \quad (7.2)$$

so that the strength of the detents was modulated by $S(p)$, which is shown in Figure 7.6, where the constants $l > 0$, $M_0 > 0$. The parameter M_0 served as the maximum detent strength to prevent the detents from becoming unreasonably strong if the musician ever pressed down especially hard.

Spring Force (SPR)

For the sake of comparison with O'Modhrain's experiment, we introduced the SPR condition, where $F_{SPR}(y)$ described a spring with stiffness k_y :

$$F_{SPR}(y) = -k_y(y - y_0). \quad (7.3)$$

The spring rest position $y_0 = -5\text{cm}$ placed the virtual spring center near the edge of the workspace. This caused the spring with relatively small stiffness $k_y = 0.07\text{N/cm}$ to exert a force to the left over the entire range of the notes included in the melodies.

No Feedback (NOFB)

We further introduced a control condition, in which the force in the y -dimension was always zero:

$$F_{NOFB}(y) = 0. \quad (7.4)$$

7.6.3 Further Ideas From Literature

We believe that we selected appropriate force conditions in the previous section, but for completeness, we mention a few other approaches mentioned in the HCI literature for assistance in selecting items from a menu.

Miller and Zeleznik recommended making detents by implementing the laws of physics describing interaction with a notch in a physical surface [109]. We tested a similar model, but we informally found the significant up and down motion corresponding to the thimble sliding into the physical notch to be distracting. The vertical motion was irrelevant to the fundamental frequency selection task at hand, so there was no need to render significant vertical motions.

Hurst et al suggested using “dirty desktops”, in which targets that are frequently chosen become more and more attractive, and less frequently chosen targets become less attractive [87]. While this method was certainly worth considering, we surmised that it would be difficult for musicians to remember how often they had chosen various notes. In that case, the haptic assistance algorithm would have been so complicated that it would have appeared non-deterministic to the musician. Thus it would not have been controllable by the musician using open-loop control.

For a micromanipulation application, some researchers recommended computing optimal paths and connecting the thimble to the optimal path using a viscoelastic link as well as creating repulsive force-fields around obstacles [5]. However, we felt that this form of haptic assistance was not applicable to the one-dimensional fundamental

frequency selection task. Firstly, there were no obstacles per se, and secondly, it was not possible to define an optimal path in general since we did not want the assistance algorithm to know which notes the musician would select at which times. A trajectory estimator could have been developed to guess what notes the musician was aiming for, but the estimator would have been subject to statistical error, so any significant use of the estimate would have resulted in a non-deterministic haptic assistance algorithm.

7.7 Subject Test

To verify the efficacy of haptic assistance in this scenario, we conducted a subject test comparing how accurately subjects could select fundamental frequencies under the DET, FRC, SPR, and NOFB force conditions. Subjects were asked to play notes from some simple melody excerpts under the different force conditions. In order to demonstrate that haptic assistance did not necessarily prevent subjects from playing smooth pitch inflections, as is characteristic to the Theremin's instrumental practice, each melody excerpt also contained a glissando.

7.7.1 Experiment Design

The experiment design was motivated by the desire to obtain data that could be compared across subjects, even if some learning effects could possibly be seen in the data. We recruited seven subjects from the Stanford Symphonic Orchestra and nine subjects from the M.S. and Ph.D. programs at the Center for Computer Research in Music and Acoustics (CCRMA) at Stanford University. Of the sixteen subjects total, one was left-handed and five were female.

The total time commitment required of each subject was about one hour on average, and each subject was compensated with \$20 for his or her efforts. Each subject was trained and tested as follows:

1. At the beginning, the subject was asked to begin filling out a questionnaire, which is given in Figure B.2.

2. The subject was presented with a sheet containing the seven simple melody excerpts shown in Figure 7.7. The first excerpt was for training, while the remaining excerpts were for testing. Each excerpt contained an average of about 9 notes from the C major scale and was followed by a glissando (see Figure 7.7). To help the subject learn the excerpts before being tested, the subject was asked to play them once on a standard piano keyboard-based instrument. Then the subject was asked to play them again to reinforce learning of the excerpts.
3. The subject was introduced to the musical instrument described in Section 7.5. The subject was asked to feel each force condition using his or her right hand, while the operation of the force condition was explained. The subject was instructed to try to use the following strategy when using the FRC condition: the subject should press down slightly when playing small intervals and lift up slightly or maintain a neutral hand weight when playing larger intervals or glissandi.
4. The process for recording performances was explained to the subject as indicated in Figure B.1. This process modeled the procedure for recording a part of a song in a music studio and consisted of the steps 1) listen, 2) practice, 3) perform, 4) consider whether to move on. In the third step, the subject performed an excerpt for a given force condition along with a metronome track. In the fourth step, the subject could elect to re-record a performance if he or she were unsatisfied with how he or she played the given excerpt.
5. Next the subject practiced using the recording process and manipulating the PHANTOM by recording his or her performance of the single training melody excerpt (see Figure 7.7, top) for each of the force conditions. Any additional questions were then answered.
6. Finally the subject recorded himself or herself performing according to the four force conditions. In order to minimize the ordering effects of the subjects' being tested on one force condition before another, the force conditions were presented across subjects according to a balanced Latin square [72]. However, during the

Training:



Testing:



Figure 7.7: Melody excerpts

testing of each force condition, the six test melodies were presented in always the same order. We chose to do so because we believed that since the subjects learned the (very simple) melodies before performing the test, the ordering of the force conditions would affect the results much more than the ordering of the melodies. Moreover, in this configuration, it was more straightforward to compare the performances of the melodies for a given force condition across subjects.

7. The subject was asked to finish filling out the questionnaire (see Section B.1). The comments that each subject wrote on his or her questionnaire are listed in Section B.2.

7.7.2 Data Analysis

We analyzed the fundamental frequency contours that the test subjects played to evaluate the tuning of their performances. In order to employ a unit approximately reflecting how humans perceive the tuning, we represented the fundamental frequency using MIDI note numbers. The MIDI note number for “middle C” is 60. Since each unit of the MIDI scale represents a half step, the first C# above middle C has the MIDI note number 61, the first D above middle C 62, and so forth. Similarly each percent of a unit represents one cent since there are 100 cents in a half step.

Let $P_s(m, c, t)$ be the fundamental frequency played by subject number s , where m indicates which melody excerpt, c indicates which force condition, and t is time in seconds. $P_7(5, c, t)$ for the NOFB condition is shown in the thin black line in Figure 7.8. For each performance, we segmented the portion of the contour describing the notes played. Segmentation was carried out by hand using a graphical user interface in MATLAB in order to compensate for timing differences in the subjects’ performances, allowing us to minimize the influence of timing differences on our measure of the accuracy of the fundamental frequencies selected. The segmented portion of the notes for $P_7(5, \text{NOFB}, t)$ is shown in a thick blue line in Figure 7.8.

To obtain an estimate of the discrete note sequence that a subject intended to play, we quantized the measured MIDI note contour to the nearest MIDI notes from




Figure 7.8: Fundamental frequency contour for subject 7 and melody excerpt 5 for the NOFB force condition (whole unsegmented contour in thin black line, segmented contour of notes played in thick blue line, and estimated intended notes shown in thick dash-dotted red line)

the C major scale. However, there was a tendency to detect some spurious notes, especially at transitions between distant notes. To greatly reduce this tendency, we eliminated notes with durations shorter than 0.2 sec, starting with the longest of these spurious notes first. As each spurious note was eliminated, it was replaced by its neighbors, where the transition time between neighboring notes was chosen to minimize the error measure described below. We entitled our quantization algorithm the Estimated Intended Note Quantizer, or EINQ. $\text{EINQ}(P_7(5, \text{NOFB}, t))$ is shown in a thick dash-dotted red line in Figure 7.8. The contour never takes on the values 68, 70, and 73 because these notes were not included in the C major scale. Note that $\text{EINQ}(P_7(5, \text{NOFB}, t))$ successfully filtered out the spurious notes that would otherwise have been detected during the thirds played at 7 sec, 8 sec, and 11 sec.

Mean Absolute Error (MAE)

To evaluate the quality of the tuning of each performance, we calculated the norm between the EINH contour and the actual performed fundamental frequency contour. We chose to use the mean absolute error (L_1 -norm of the error) instead of the root mean square (L_2 -norm of the error) so that the error measure did not emphasize the inevitably large error contributions stemming from note transitions, but rather focused more on the typically more constant error contributions from within notes (see the contours in Figure 7.8). The mean absolute error (MAE) of the fundamental frequency contour, from time T_1 to time T_2

$$MAE_s(m, c) = \frac{1}{T_2 - T_1} \int_{T_1}^{T_2} |P_s(m, c, t) - \text{EINH}(P_s(m, c, t))| dt. \quad (7.5)$$

T_1 and T_2 denote the segmentation time boundaries of the notes played prior to the glissando during the melody excerpt. $T_1 = 4\text{s}$ and $T_2 = 11.2\text{s}$ in Figure 7.8 as highlighted by the thick blue and dash-dotted red lines.

RMS Acceleration

We wanted to study to what extent detents caused the musician to play jagged fundamental frequency inflections rather than smooth fundamental frequency inflections. To this end, we manually segmented the portion of each performance corresponding to a glissando. For example, in Figure 7.9 (top) the thick blue line shows the segmented portion of the performance $P_7(5, NOFB, t)$ corresponding to the thick dash-dotted black line. We evaluated the jaggedness with the root-mean-squared (RMS) acceleration of the fundamental frequency trajectory $\ddot{P}_s(m, c, t)$ during each glissando, where $[T_3 \ T_4]$ denotes the segmentation time interval of each glissando:

$$Accel_s(m, c) = \sqrt{\int_{T_3}^{T_4} \frac{\ddot{P}_s^2(m, c, t)}{T_4 - T_3} dt}. \quad (7.6)$$

To estimate the fundamental frequency acceleration $\ddot{P}_s(m, c, t)$, we filtered the fundamental frequency trajectory $P_s(m, c, t)$ both forward and backward by the following




Figure 7.9: RMS acceleration of fundamental frequency contours for subject 7 and melody excerpt 5 for the NOFB force condition (TOP: unsegmented fundamental frequency contour in thin, dash-dotted black line; segmented contour in thick, blue line; BOTTOM: unsegmented acceleration of fundamental frequency contour in thin, dash-dotted black line; segmented contour in red)

filter:

$$H_{DIFFERENTIATOR}(z) = (1 - p) \cdot \frac{1 - z^{-1}}{1 - pz^{-1}} f_s. \quad (7.7)$$

We filtered forward and then backward in time so that the acceleration estimate would not be delayed in time. The sampling rate was $f_s = 2.4\text{kHz}$, and we chose $p = 0.995$ so that the higher frequencies would not be emphasized so much as to contribute significant noise to the acceleration estimate. The estimated fundamental frequency acceleration $\ddot{P}_s(m, c, t)$ is shown in dash-dotted black in Figure 7.9 (bottom), with the segmented portion shown in solid red (see Figure 7.9, bottom).




Figure 7.10: Basic detent element

7.7.3 Selecting The Parameters For DET

The basic detent element is shown in Figure 7.10. Only limited information was available on tuning the precise shape of the basic detent element. For example, researchers reported finding in informal tests that altering the shape of the detent nonlinearly manifested itself psychophysically as apparently only an intensity difference [173][125]. Consequently, we decided to devise a small test to help us fine tune the shape of our basic piecewise linear detent element, on which the DET and FRC conditions depended. The basic detent element is described mathematically for any horizontal position $y \in [-W \ W]$:

$$F_{DET}(y) = -\text{sign}(y) \cdot \begin{cases} \frac{|y|D}{\rho W} & \text{if } |y| < \rho W \\ \frac{(W-|y|)D}{(1-\rho)W} & \text{otherwise} \end{cases} \quad (7.8)$$

W described the radius of the detent, D described the maximum absolute value of the force, and $\rho \in (0 \ 1)$ adjusted the y -position corresponding to the maximum force level (see Figure 7.10). Figure 7.2 can serve as a reminder to the reader about the orientation of the y -axis. $W = 0.375\text{cm}$ was already fixed due to the approximate




Figure 7.11: Fundamental frequency contour played by one of the authors subject to the DET condition for $D = 0.5N$ and $\rho = 0.3$

20cm workspace and the desire to allow play over a range of about two octaves. The remainder of this section explains how we chose ρ and D .

We tested one of the authors playing an ascending diatonic scale followed by a descending glissando. The fundamental frequency contour is shown in Figure 7.11 for $D = 0.5N$ and $\rho = 0.3$. Since the detents were fairly strong, they caused the glissando between 7.5sec and 9.5sec to be a little bit bumpy or jagged.

Sweep 1

In order to help motivate our choices of D and ρ , we tested the author for various choices of D and ρ while playing the ascending scale followed by the descending glissando. Here we present the results for two sweeps. For the first sweep, ρ was held constant at 0.3 and D was swept from almost zero to about 1.2N. The MAE for the diatonic scale is shown in Figure 7.12. Increasing D corresponds to increasing the degree of haptic assistance, so as expected we observed that increasing the haptic assistance reduces the MAE of the fundamental frequency contour played during the diatonic scale. There is a knee in the curve near $D = 0.6$ (see Figure 7.12).




Figure 7.12: Sweep 1: MAE for one of the authors playing the diatonic scale with $\rho = 0.3$ while D was swept from about zero to about 1.2N

We plot the RMS acceleration as a function of D in Figure 7.13. Clearly the RMS acceleration increases as the degree of haptic assistance is increased. In other words, the glissando becomes more jagged or bumpy for the DET condition as the degree of assistance is increased. This plot illustrates why we do not want to choose D to be too large.

Sweep 2

In a second sweep, we held D constant at 0.41N, while we varied ρ from small values to nearly 0.5. The MAE of the notes played and the RMS acceleration of the glissando for this sweep are shown in Figures 7.14 and 7.15, respectively. In general, ρ appears to have had little effect on the quantities we use for analysis. However, it is clear that for very small ρ , the MAE of the notes decreased slightly, presumably due to the steep slope of the detent locally about its center. In other words, reducing ρ caused the MAE to decrease slightly because the haptic assistance was more focused. However, if ρ was chosen to be much smaller than 0.15, we found informally that it was easy to get stuck in unintended detents. In other words, the haptic assistance was too focused and thus distracting.

We performed a few tests to optimize the basic detent element for our DET




Figure 7.13: Sweep 1: RMS acceleration for one of the authors playing a glissando with $\rho = 0.3$ while D was swept from about zero to about 1.2N




Figure 7.14: Sweep 2: MAE for one of the authors playing the diatonic scale with $D = 0.41\text{N}$ while ρ was swept from nearly zero to 0.5




Figure 7.15: Sweep 2: RMS acceleration for one of the authors playing a glissando with $D = 0.41\text{N}$ while ρ was swept from nearly zero to 0.5

condition. The tests explained above guided us to choose $D = 0.41\text{N}$ and $\rho = 0.3$. There was a trade-off in choosing D . If D was too large, then the haptic assistance was so strong that it prevented the musician from being able to play smooth glissandi with DET (see Figure 7.15). On the other hand, if D was too small, then the haptic assistance did not reduce the MAE as much as possible (see Figure 7.14).

Choosing the optimal ρ parameter appeared to be less crucial. Perhaps this was because the potential energy change E_{DET} due to moving the thimble through one side of the detent is independent of ρ :

$$E_{DET} = \int_0^W F_{DET}(y)dy = -\frac{WD}{2}. \quad (7.9)$$

Nevertheless, according to informal testing, altering ρ changed the way the detents “felt”—in some sense, it changed their texture.

During pilot tests with the melody excerpts themselves, we decided to make the parameters slightly more focused by changing the parameters to $D = 0.41\text{N}$ and $\rho = 0.2$. These slightly more aggressive parameters worked well when tested by the author for the melody excerpts shown in Figure 7.7. We had the impression that reducing ρ may have reduced the tendency to get stuck in unintended detents, but we made no attempts at verifying this formally. Finally, we present the finalized shape




Figure 7.16: Finalized basic detent element shape with $\rho = 0.2$, $W = 0.375\text{cm}$, and $D=0.41\text{N}$

of our basic detent element where $\rho = 0.2$, $W = 0.375\text{cm}$, and $D=0.41\text{N}$ in Figure 7.16.

7.7.4 Qualitative Results

We begin by looking at the fundamental frequency trajectories for subject 7 and melody 5. These particular trajectories were chosen because they provide good examples for illustrating the typical behavior of the test subjects. Figure 7.17 shows the trajectory for the NOFB condition, where no haptic feedback was provided. The subject's trajectory is shown in blue, while the estimated intended note quantization (EINQ) is shown in dash-dotted red. Finally, the glissando portion is indicated by the especially thick green line. The subject was able to perform the melody; however, the performed and intended notes do not match up well together. As a consequence, $MAE_7(5, NOFB) = 0.43$. In other words the mean absolute fundamental frequency error was about 0.43 half steps, or 43 cents. The performance was similar in the case of spring (SPR) feedback (see Figure 7.18). We believe that the NOFB and SPR force conditions resulted in large fundamental frequency MAE because they provided




Figure 7.17: NOFB No Haptic Feedback (subject 7, melody 5, $MAE=0.43$, $Accel=9$)

no physical indication guiding to the correct fundamental frequency locations.

In contrast, the multiple detents (DET) condition caused the performed fundamental frequency trajectory to match the intended note sequence much better as evidenced by the decreased $MAE_7(5, DET) = 0.23$ (see Figure 7.19). However, the detents were strong enough that the glissando was somewhat jagged or bumpy (see the thick, green line in Figure 7.19). The elevation of the RMS acceleration $Accel_7(5, DET)=16$ indicated the same.

The performance given the force-sensitive detents (FRC) is shown in Figure 7.20. The test subject modulated the strength of the detents appropriately. The notes were selected accurately with the help of haptic assistance, while the glissando was smooth because the test subject was able to disable the haptic assistance when desired.




Figure 7.18: SPR Spring Feedback (subject 7, melody 5, $MAE=0.43$, $Accel=14$)




Figure 7.19: DET Multiple detents (subject 7, melody 5, $MAE=0.23$, $Accel=16$)




Figure 7.20: FRC Force-Sensitive Detents (subject 7, melody 5, $MAE=0.17$, $Accel=10$)

7.7.5 Results

Basic Statistical Analysis

We performed a statistical analysis on the subject test data in order to verify the significance of the findings. Figure 7.21 shows a box plot of the mean absolute error (MAE) in selecting the fundamental frequencies of notes subject to the four different force conditions. The whiskers are the dashed lines that show the extent of the data, while the red plus signs outside of the extent represent outliers. We made no effort to remove outliers, which were most likely due to subject error, from the data set because our results were already statistically significant. Each blue region reaches from the lower quartile to the upper quartile of the data. The red line within each blue region shows the median of the data.

The notches in the box plots are important. If the notches for two conditions do not overlap, then the data for the conditions are drawn from two different distributions with likelihood greater than 95% [107]. In other words, if the notches for




Figure 7.21: Box plot of the mean absolute error (MAE) of note fundamental frequency trajectories across all subjects and melody excerpts

two conditions do not overlap, then the difference between the performance given the two conditions is considered statistically significant. Figure 7.21 clearly shows that the FRC and DET force conditions caused the fundamental frequency MAE to decrease at a statistically significant level. In addition, FRC appears to have helped the test subjects select fundamental frequencies slightly more accurately at a marginally (95%) significant level.

The analogous analysis of the RMS acceleration of the fundamental frequency trajectories of glissandi is shown in Figure 7.22. Ten outliers were present in the data spread across the force conditions, but the outliers were cropped away from the box plot to aid the reader in focusing on the remaining data. As expected, the RMS acceleration for the multiple detents (DET) condition is significantly elevated in comparison with the other force conditions. The basic detents, which helped the test subjects select the fundamental frequencies of notes accurately, interfered with playing pitch inflections such as glissandi (see also Section 7.7.3). However, Figure 7.22 shows that the force-sensitive detents (FRC) condition resulted in a statistically significant




Figure 7.22: Box plot of the RMS acceleration of the fundamental frequency trajectories of glissandi across all subjects and melody excerpts

lower RMS acceleration, which was comparable with the no feedback (NOFB) and spring feedback (SPR) conditions. Since the test subjects were able to modulate the strength of the detents for the FRC condition, the haptic assistance did not cause them to play excessively jagged or bumpy glissandi.

Due to our desire to avoid fatiguing the test subjects, the training session was brief (see Section 7.7.1). As a consequence, a few of the subjects did not learn to modulate the degree of haptic assistance by changing the downward force when using the FRC condition. While only one subject admitted on the questionnaire that she did not “apply more downward force when playing small intervals and less downward force when playing larger intervals or glissandi”, when we look at the FRC fundamental frequency trajectory data, we estimate that about four of the subjects did not apply little enough downward force while playing the glissandi. As a consequence, we believe that if we had provided the subjects with more training time, they would have been able to better leverage the haptic assistance. In particular, we believe that the FRC condition would have provided even better results since it was the most complicated




Figure 7.23: Preferences of the test subjects

to learn to use. Hence, we believe that the mean for FRC shown in Figure 7.22 would have been further decreased.

Questionnaire

The sixteen subjects were asked to order the force feedback conditions from the most to the least preferable for playing glissandi accurately as well as for playing notes accurately. Fifteen subjects most preferred either FRC or DET for playing notes accurately. Of these subjects, eleven most preferred FRC for playing notes accurately.

Most subjects preferred the non-detent conditions for playing glissandi accurately: eleven subjects most preferred either NOFB or SPR for playing glissandi accurately.

However, the remaining five subjects most preferred the FRC condition for playing glissandi accurately, presumably because these subjects had learned sufficiently well to lift the thimble up slightly while playing glissandi, so as not to get distracted by the detent-based haptic assistance.

Residual Learning During Test Trials

We wanted subjects to learn to play the melody excerpts on a piano keyboard-based instrument sufficiently well prior to the test trials. However, to avoid fatiguing the subjects, we only had time to ask each subject to play the excerpt twice on the piano keyboard-based instrument prior to testing (see Section 7.7.1). Accordingly, we deliberately selected the easy-to-learn melody excerpts shown in Figure 7.7 for testing.

Nevertheless, during the test trials, subjects inevitably learned the melody excerpts better, so we performed further analysis to verify that this effect was minimal. Figure 7.24 shows the mean absolute error (MAE) of the fundamental frequency trajectory plotted as a function of time, measured using the test trial time index. Test trial time indices were used since subjects received the melody-condition pairs in four different orders. Index 1 was the first melody-condition pair that a subject was tested on, and index 24 was the last melody-condition pair that a subject was tested on. The data was plotted across all test subjects (see Figure 7.24, blue dots). We fit an affine regression model to the data. The red line in Figure 7.24 shows the model, which fit the data as closely as possible in a least-squares sense. The affine regression model suggests that subjects improved by 0.01 half steps (i.e. 1 cent) on average over the course of the experiment. This learning was, however, dwarfed by the much larger task MAE. Averaging over the entire dataset (all test subjects, all melodies, and all force conditions), the average MAE was 0.33 half steps (or 33 cents). We conclude that the subjects learned the melodies well enough prior to the test trials that residual learning had only a minor effect on the data, which explains how we were able to obtain useful and statistically significant results from the experiment.




Figure 7.24: Fundamental Frequency MAE shown as a function of time during the test trials to measure residual learning during testing

Melody Excerpt Difficulty

Some melody excerpts were more difficult for the subjects to play accurately than others. For reference, we present a box plot of the fundamental frequency MAE in Figure 7.25, which shows that while test melodies one and two were the easiest, test melodies four and six were particularly difficult. Figure 7.7 shows the melody excerpts themselves.

Conclusions

Haptic assistance for musicians will open up a whole range of new instruments that can aid musicians in playing. In the framework of a formal subject test, we demonstrated that haptic detents (DET and FRC) improve the accuracy with which musicians select fundamental frequencies over a linear range in comparison with a simple spring force (SPR) or no haptic feedback at all (NOFB). Force-sensitive detents (FRC) are superior to basic detents (DET) due to the musician's ability to regulate the amount of haptic assistance in real time. FRC allows musicians to press downward more when more haptic assistance is required, for instance when honing in on a specific note fundamental frequency or when playing a small interval, and FRC allows musicians




Figure 7.25: Box plot of fundamental frequency MAE across the test melody excerpts

to press downward less when less haptic assistance is required, such as when playing large intervals or glissandi. The FRC condition is so helpful that we believe it will be applied in other areas of haptics, such as for assisting users in selecting items from menus.

In contrast with prior work [120], we did not observe any statistically significant difference between no haptic feedback and the very simple spring haptic feedback (SPR). However, the test subjects did prefer the SPR condition to the NOFB condition. Perhaps if we had tested an order of magnitude larger number of test subjects, we would have been able to measure a statistically significant difference between NOFB and SPR.

7.8 Final Words

At first our conclusions might seem to be at odds with O’Modhrain’s experiment. Indeed while O’Modhrain suggested simply that some form of haptic feedback was better than none at all [120], we suggest that detents are even more helpful in assisting musicians in selecting fundamental frequencies over a modest linear range. However, we think that the real lesson to learn is that *the best type of haptic assistance depends*

on the task. In fact, this was the motivation behind the development of the novel force-sensitive FRC condition, where the degree of haptic assistance depends directly on user input. We believe that future haptic assistance devices should consider taking advantage of real-time user input in a similar manner, allowing users to constantly regulate the assistance that they receive.

Chapter 8

Conclusions

8.1 Review

By incorporating haptic feedback into musical instruments, designers can facilitate rich physical interactions between a musician and a musical instrument. In Appendix A, we introduce a general method for designing haptic musical instruments that allows two-directional mechanical interaction between a musician and a virtual acoustic musical instrument. However, haptic feedback has many more important applications. In the both rhythm and pitch contexts, we demonstrated that haptic feedback can be used to assist a musician in making gestures more accurately. In a rhythmic context, we demonstrated further that haptic feedback could assist a musician in making gestures that would otherwise be difficult or impossible. These are some of the most compelling reasons why musical instrument designers should consider incorporating active haptic feedback into their musical instruments.

8.2 Collocation

In order for haptic feedback control to exercise a significant amount of control authority, sensors and actuators should be collocated. Perfect collocation implies that if a positive real or any other passive controller is applied to the control system, then the system is unconditionally stable (see Section 2.4); in other words, the loop gain

can be made arbitrarily large without destabilizing the system. For more details on the challenges of designing a perfectly collocated control system, see Section 2.5.

8.3 Physical Analog Models

Physical analog models can be used not only for simply understanding the behavior of dynamical systems, but also for designing new feedback controllers. An equivalent physical analog model exists for each positive real controller and provides a metaphor by which to understand the effect of the controller. Consequently, positive real controllers are useful for designing both Feedback Controlled Acoustic Musical Instruments and Haptic Musical Instruments (see Chapter 3). Many useful nonlinear controllers correspond to nonlinear physical analog models as well. For these reasons, it is important that musical instrument designers understand physical analog models and how to apply them.

Humans in general are also intuitively familiar with physical systems due to day-to-day interactions. As humans interact with the world, they create conceptual models for interacting with physical systems [169]. Humans can combine multiple models together to synthesize more complex models [139]. In addition, humans can transfer skills from one system to another if the dynamics are similar [86]. For these reasons, we argue that humans should be able to quickly learn to interact with new musical instruments whose designs conform to human physical intuition. We believe that this is one of the reasons why Physical Interaction Design for musical instruments is becoming so popular [165][111].

8.4 Further Advantages of Active Feedback

Active feedback systems can emulate passive feedback systems by implementing passive controllers. In some scenarios, it would be useful to have a system that is passive but whose impedance varies over time. Often it is much easier to implement such systems with active feedback where a digital feedback controller changes the control parameters over time. For instance, it is easy to make a active feedback system

emulate a time-varying mass, damper, or spring.

In addition, in contrast with passive systems in nature, active feedback systems are capable of adding energy to the musical instrument interface. This characteristic can be useful in many circumstances. The Haptic Drumstick and Haptic Drum assist the musician in playing drum rolls by actively promoting limit cycle-type behavior (see Chapter 6). In addition, the FRC condition, the most effective force feedback condition for assisting musicians in selecting fundamental frequencies as described in Chapter 7, is also active. In particular, it is an active feature that the strength of the horizontal force in the y -direction can be modulated by adjusting the force in the z -dimension. It is possible to prove this by connecting a spring to the haptic device in a manner that causes the system to self-oscillate according to a stable limit cycle. Active feedback systems will continue to be employed in situations where a dynamical system with an internal energy source is required.

8.5 Access to Feedback Technology

In the past, force feedback technology has been less accessible to musicians due to elaborate software requirements and the expense of quality devices providing collocated sensing and actuation. In Appendix C, we provide a new toolbox that facilitates rapid prototyping of Haptic Musical Instruments using the inexpensive Novint Falcon device. Similarly, the Toolbox for the Feedback Control of Sound, as described in Appendix G, makes it more convenient for musicians to construct Feedback Controlled Musical Instruments.

8.6 Haptic Assistance Guidelines

Finally, we provide some guidelines for designing assistive haptic interfaces. These guidelines are based on our experiences in designing the haptic musical instruments outlined in chapters 5-7. For additional guidelines on creating haptic systems in general, see the paper by Carter and Fourney. Most notably, many of their guidelines are based on the notion that users can interact intuitively with haptic systems that

correspond to physical analog models [44].

1. The most effective type of haptic assistance depends on the task (see Section 7.8).
2. It can be helpful to allow the user to regulate the degree of haptic assistance in real time (see Section 7.7.5).
3. If haptic assistance is deterministic and relatively simple, then it is more likely to be useful (see Section 5.4.6).
4. Effective haptic assistance should never take an unexpected action so fast that it surprises the user, which could prevent him or her from responding appropriately (see Section 5.4.6).
5. Many useful haptic assistance systems are not passive (see Section 8.4).

Appendix A

Using Haptic Devices to Interface Directly with Digital Waveguide-Based Musical Instruments

A.1 Introduction

Other researchers have previously designed haptic musical instruments that mimic real acoustic musical instruments. The feedback controllers have been implemented using finite difference and (approximate) hybrid digital waveguide models. We present a novel method for constructing haptic musical instruments in which a haptic device is directly interfaced with a conventional digital waveguide model by way of a junction element, improving the quality of the musician's interaction with the virtual instrument [22]. We introduce both the explicit digital waveguide control junction and the implicit digital waveguide control junction.

A.1.1 Prior Work

Many haptic musical instruments consist of a haptic device coupled to a virtual musical instrument, which is often implemented by way of a physical model. The common resonant physical model types in the literature with applications to haptics are finite difference, finite element, and digital waveguide.

A.1.2 Finite Difference Modeling

Haptic musical instruments have been developed since the 1970's at the latest. The Association pour la Création et la Recherche sur les Outils d'Expression (ACROE) currently has the longest history of model development in physical modeling for haptics [41]. The Cordis Anima system allows virtual masses to be connected to one another by way of viscoelastic links emulating springs and dampers [155]. Jean-Loup Florens and François Poyer explain how they use nonlinear viscoelastic links to allow virtual objects to be bowed by way of a haptic device [63][129]. The differential equations describing the masses and links are discretized and solved approximately using finite difference techniques [155]. ACROE's Genesis environment aids in synthesizing and editing large-scale models [45].

A.1.3 Finite Element Modeling

Finite element models have been computed in real-time for applications in haptics, but the computational requirements can be prohibitive [43]. The computation can be reduced somewhat if the model is static [36]—i.e. if no state is required to keep track of the history of the model's motion; however, these models are not useful for designing musical instruments that exhibit resonance behavior.

Along similar lines, Cynthia Bruyns used a finite element model to compute modal synthesis parameters [37]. Her model is more computationally efficient and could probably be applied successfully in haptics applications.

A.1.4 Digital Waveguide Modeling

In contrast with other modeling methods, digital waveguide modeling is especially efficient because of the small number of multiply-adds and the large portions of the memory that can be factored into digital delay lines. Furthermore, when modeling most one-dimensional waveguide-based instruments, digital waveguide modeling is equivalent to discretizing the solution to the wave equation rather than approximately solving the wave equation [149]. Consequently, such digital waveguide models are more accurate than finite element and finite difference models. Finally, there is an extensive literature on calibrating digital waveguide models so that they sound like real musical instruments [149].

However, despite the advantages of digital waveguides, their application in haptics has been quite limited. Rather than directly coupling a model to a haptic device, previous research has focused on more complicated hybrid digital waveguide models. Consider the structure shown in Figure A.1. A grossly simplified haptic model, such as a nonlinear spring, is typically used to govern the dynamics of the haptic device [141]. The musician's physical interaction with the instrument is limited by the simplification. Nevertheless, sound may still be synthesized as the parameters describing the state of the haptic device are fed directly into a digital waveguide model (see Figure A.1). Plucked strings¹ [60][64], a bowed string [141], and a drum [27] have been previously implemented using this approach. Kontogeorgakopoulos has studied a related problem in which he simulates interfacing digital waveguides with Cordis Anima finite element networks [100].

A.2 Direct Modeling Approach

To improve upon the hybrid digital waveguide modeling approach, while still retaining the advantages of digital waveguides, we propose using a control junction element to directly connect a haptic device to a digital waveguide according to the signal flow diagram shown in Figure A.2. We introduce both explicit and implicit variations of

¹Rob Shaw developed a haptically plucked string at Interval Research Corporation using the hybrid modeling approach, but as far as we know, he did not publish his research results.




Figure A.1: Signal flow diagram for hybrid modeling approach

the junction.

A.2.1 Explicit Control Junction Element

The explicit control junction element for a velocity wave model is shown inside the dotted box in Figure A.3. The junction consists of a sensor and an actuator placed only one sampling interval apart along a digital waveguide with wave impedance R_0 . For a slightly more general digital waveguide model, see Figure 2.4. The signal a flowing to the right along the top models the velocity wave traveling to the right, while the signal b flowing to the left along the bottom models the velocity wave traveling to the left. The velocity output $v_s[n]$ at the control point is the sum of the two traveling wave components at that point a_L and b_L . The force input $F_s[n]$ causes traveling waves to exit the junction to both the left and the right. The waves are scaled by $2R_0$ since exerting a force on the control point of a virtual string should be analogous to exerting a force on the ends of two freely terminated virtual strings in parallel, each with wave impedance R_0 [149].

We elucidate the operation of the control junction with the help of an example




Figure A.2: Signal flow diagram for direct modeling approach

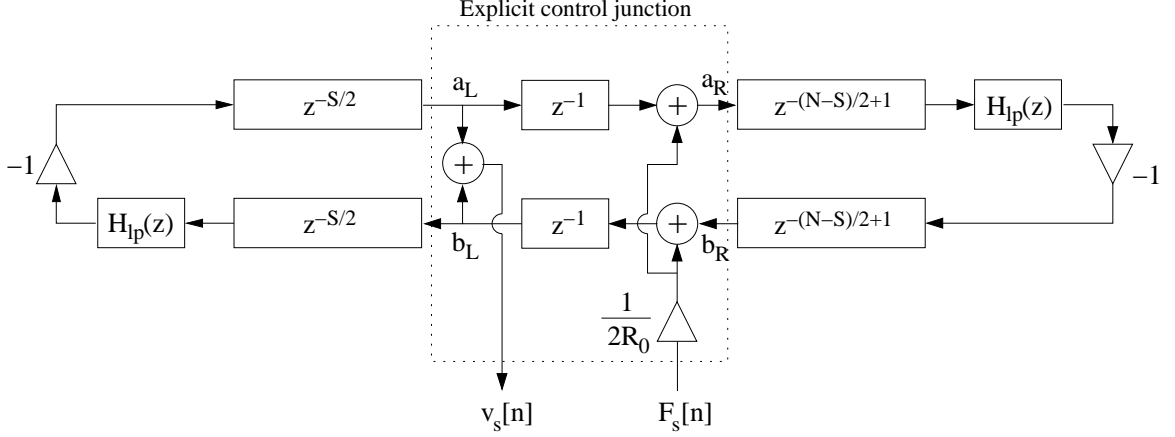


Figure A.3: Explicit control junction connected to a digital waveguide model of a vibrating string

vibrating string model (see Figure A.3 in its entirety). The top three delay lines $z^{-S/2}$, z^{-1} , and $z^{-(N-S)/2+1}$ model velocity waves traveling to the right, and the bottom three delay lines $z^{-(N-S)/2+1}$, z^{-1} , and $z^{-S/2}$ model velocity waves traveling to the left. S/N is proportional to the distance from the left string end to the control point divided by the total string length. If f_s is the sampling rate, then f_s/N is approximately equal to the string's fundamental frequency. $H_{lp}(z)$ is a linear phase low-pass filter that causes higher string resonances to decay more quickly. $H_{lp}(z)$ further promotes stability by reducing the gain of high frequency energy reflected from the virtual string terminations.

A.2.2 Interfacing With The Waveguide

We can now apply standard teleoperation techniques to interface a haptic device with the control junction. It is tempting to try to bind the haptic device and virtual instrument directly together; however, this approach can be tricky because the haptic device will be able to render only a finite maximum stiffness [56]. If the instrument model were to specify a greater stiffness, then the system could become unstable. Teleoperator theory suggests instead that we should bind the haptic device and the virtual model together with a spring that limits the maximum stiffness that the haptic




Figure A.4: Teleoperator control of a virtual vibrating string

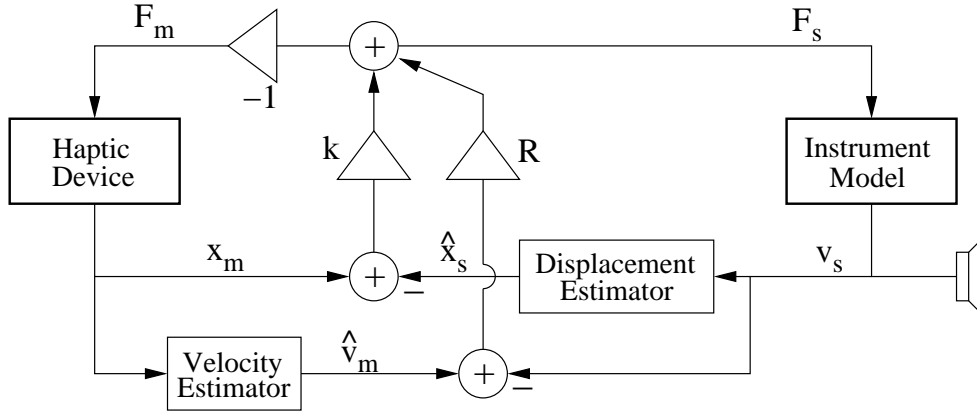


Figure A.5: Signal flow diagram detailing how to implement teleoperator control

device needs to render. Furthermore, adding a damper as well, as shown in Figure A.4, can help eliminate tendencies for the system to become unstable.

Some additional details must be considered. For instance, many haptic devices measure displacement, while the instrument model provides only a velocity output. It follows that some estimators will be required for our example. Figure A.5 shows how a velocity estimator can be employed to estimate the haptic device velocity $\hat{v}_m[n]$ from the haptic device displacement $x_m[n]$. Similarly, a displacement estimator estimates the displacement of the vibrating string $\hat{x}_s[n]$ from the string velocity $v_s[n]$ (see Figure A.5).

The force due to the spring k and damper R is thus:

$$F_s[n] = k(x_m[n] - \hat{x}_s[n]) + R(\hat{v}_m[n] - v_s[n]) = -F_m[n]. \quad (\text{A.1})$$




Figure A.6: Helmholtz motion of haptically bowed virtual string

By allowing k and R to vary over time, we open up possibilities for simulating a wide variety of systems. To simply bind the haptic device and the virtual control point together, we use a traditional spring and damper as shown in Figure A.4 [56]; to implement plucking, we neglect the damper while implementing a spring that disengages at relatively large force levels [60]; and to implement bowing, we eliminate the spring and implement a nonlinear damper [63]. Figure A.6 shows the Helmholtz motion of the virtual string when bowed in the laboratory on a Phantom Model T haptic device operating at a servo/audio sampling rate rate of 19kHz.

The two single sample delay units in the explicit control junction (see inside the dotted box in Figure A.3) allow instrument designers to connect arbitrary impedances to the waveguide with ease. However, the single sample delay units are non-physical and could conceivably cause some stability problems at high frequencies for large loop gains. Various other factors limit maximum loop gains such as amplifier, motor, and sensor bandwidth and nonlinearity, as well as the servo rate, etc.

A.2.3 Implicit Control Junction Element

The single sample delay units may be removed entirely by making use of knowledge about the impedance of the load connected to the waveguide. We solve the problem for teleoperation with a spring and a damper as shown in Figure A.4. If we eliminate

the single-sample delay blocks shown inside the dotted box in Figure A.3, then we obtain the following:

$$v_s[n] = a_L[n] + b_R[n] + \frac{F_s[n]}{2R_0}. \quad (\text{A.2})$$

To make use of substituting (A.2) into (A.1), we need to specify the displacement estimator's form, so we use a leaky integrator with pole $p \approx 1$ but $p < 1$:

$$\hat{x}_s[n] = p\hat{x}_s[n-1] + \frac{1}{f_s}v_s[n], \quad (\text{A.3})$$

and we solve to arrive at the implicit control junction element's analog to (A.1):

$$\begin{aligned} F_s[n] = \frac{2R_0}{2R_0 + R + k/f_s} & \left(k(x_m[n] - p\hat{x}_s[n-1]) \right. \\ & \left. + R\hat{v}_m[n] - (R + k/f_s)(a_L[n] + b_R[n]) \right). \end{aligned} \quad (\text{A.4})$$

The way in which the implicit control junction is interfaced with the model is shown in Figure A.7. The haptic device (not shown) is connected to the signals below $-F_s[n]$ and $\hat{v}_m[n]$ (see Figure A.7). Note how in comparison with Figure A.3, the signal flow is more complicated; however, the unit delays z^{-1} are eliminated.

Since R and k can be changed in real-time, we can use the same strategies for implementing binding and plucking as in Section A.2.2. However, (A.4) must be re-derived in the case of bowing to take the nonlinearity of the damper R into account. In a practical context, the solution results in calculating a nonlinear lookup “bowing table” to describe the effect of the nonlinear damper [149]. In fact, since the writing of our original conference paper, Sinclair et al. have published a paper to this effect [145].

A.3 Final Words

Digital waveguides may be used advantageously when designing haptic musical instruments due to their accuracy, efficiency, and ease of calibration. While we have shown one example implementation, many more haptic musical instruments can be derived along similar lines by interfacing haptic devices with one-dimensional or even

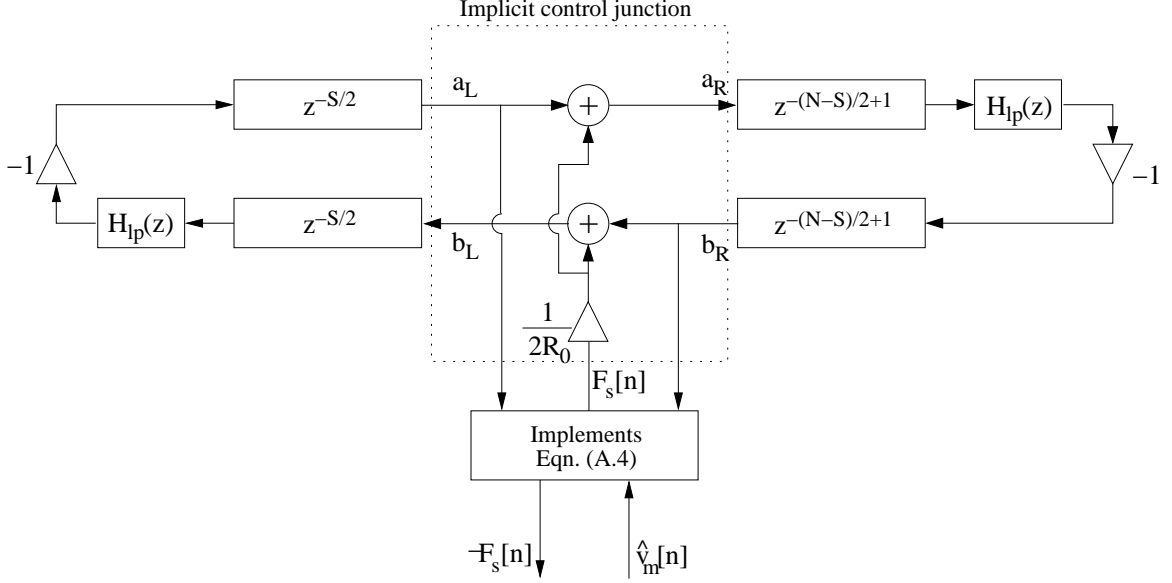


Figure A.7: Implicit control junction connected to a digital waveguide model of a vibrating string

multi-dimensional waveguide models. The explicit control junction is more straightforward to implement than the implicit control junction; however, the implicit control junction is slightly more accurate due to the absence of the non-physical delays z^{-1} (see Figure A.3).

Both the explicit and implicit control junctions conform to what we call the direct approach (as opposed to the hybrid approach), allowing the digital waveguide model to not only serve as a sound synthesis model but also to serve as a haptic interaction model, allowing the musician to interact more intimately with the instrument. For example, Figure A.8 shows the result of a musician haptically plucking a digital waveguide string (at 40ms) and then damping it (at 800ms). The string is modeled explicitly as explained in Section A.2.2. The musician's finger itself provides the damping, so as the musician changes his or her grip of the haptic device, the quality of the damping will be modified. Furthermore, especially at low fundamental frequencies such as at 50Hz, the string's displacement is large enough to interact significantly nonlinearly with the musician's finger during the damping process, resulting in a




Figure A.8: Plucking and damping using the direct approach

buzzing sound. To listen to sound examples, see the project website:

<http://ccrma.stanford.edu/~eberdahl/Projects/HDWG>

Appendix B

Further Information On Subject Test

B.1 Documents

Here we present the documents given to each of the participants during the subject test outlined in Chapter 7. The subject test was conducted according to Human Subjects Protocol 16056 “Study of Haptic Interactions in Virtual Musical Environments”, which was approved by Stanford University’s Institutional Review Board on February 27, 2009.

B.2 Questionnaire Comments

First we list the comments noted by the participants from the subject test explained in Chapter 7.

- SPR was pleasant for melody excerpt 6. Didn’t like DET for gliss. Took a little to adjust to alternating between exerting/not exerting force w/ FRC.
- Melody #6 is harder.
- I tended to like no feedback when playing normal notes I think because I play violin and am used to the “sliding” ability.

Explanation

You will be asked to perform melodies on a musical instrument implemented with a Phantom robotic arm, which can provide programmable force feedback. You will perform the set of melodies subject to four different types of force feedback. A voice will make an announcement every time the force feedback condition is being changed. The four force feedback conditions are the following:

DET: Detents or notches that provide gentle guidance toward correct pitches on the diatonic scale.

FRC: Similar to **DET** except that the strength of the guidance depends on how hard you push downward. **NOTE:** *You should push down slightly harder to receive more assistance when playing small intervals, and you should pull up slightly to receive less assistance when playing large intervals or glissandi.*

SPR: Like attaching a very loose spring to the end of the robotic arm.

NOFB: No force feedback at all.

The volume of the instrument can be adjusted *only* by changing the output volume on the sound interface. The following explains what steps you will go through to perform for each melody and force feedback condition. The steps are designed to emulate the process of recording one part for a song in the studio. You can navigate forward through the steps using the CONTINUE button.

- 1) *Listen:* You hear the melody as it is played on a French horn-like instrument.
- 2) *Practice:* You practice the melody as long as you wish in the absence of a metronome. To help you find the first note of the melody, you hear this pitch panned far to the right side. If you wish to hear the melody again, press GO BACK to return to step 1. Otherwise, press CONTINUE to move on to step 3.
- 3) *Perform:* The sound in the right ear cup ceases, allowing you to focus on your own playing, and a metronome marks time. *You begin performing after the first four beats.* That is, you should begin playing the melody as the metronome sounds the fifth time.
- 4) *Consider Whether To Move On:* The metronome stops. If you are satisfied with your performance, press CONTINUE to move on to the next melody. Otherwise, press GO BACK to return to step 2.

Try to play each melody as well as you can for the given force condition—you may re-record your performance for each case repeatedly until you are satisfied your performance.

Remember that for the **FRC** condition, *you should pull up slightly to receive less assistance when playing large intervals or glissandi.*

Figure B.1: Explanation provided to test subjects prior to the procedure

ID # _____

Please answer the following questions:

Date of birth: _____

Gender: _____

Preferred writing hand: _____

Can you sight sing? _____

Which instruments do you play? _____

How many years of musical training do you have? _____

About how many hours a week do you play any musical instruments? _____

Please answer these questions after you have completed the tasks with the haptic device:

Glissandi:

Which force feedback condition did you most prefer for playing glissandi accurately? _____

Which force feedback condition did you least prefer for playing glissandi accurately? _____

Please order the force feedback conditions from the most preferable to the least preferable for playing glissandi accurately _____

Notes:

Which force feedback condition did you most prefer for playing notes accurately? _____

Which force feedback condition did you least prefer for playing notes accurately? _____

Please order the force feedback conditions from the most preferable to the least preferable for playing notes accurately _____

When playing subject to the **FRC** condition, did you apply more downward force when playing small intervals and less downward force when playing larger intervals or glissandi?

Circle one: Yes No

Other comments _____

Figure B.2: Questionnaire

- Melodies with larger intervals were easier to play with NOFB with muscle memory rather than DET or FRC.
- The spring feedback felt most natural (more like playing an instrument) although the detents allowed better pitch accuracy.
- Overall FRC and NOFB are best in my opinion.
- For some reason, I preferred SPR to NOFB because it felt more comfortable while jumping over larger intervals and also glissandi. DET is nice for small intervals. FRC naturally felt the most comfortable with the option to do both.
- Very cool. When can I buy the Phantom Theremin Thing at Guitar Center?
- The (tuning) background noise gets annoying and confusing after a while. I would have preferred to find the pitch by myself. Thimble too large for finger.
- Great Instrument!
- Tried to keep the downward force constant to the preferred friction to move across (i.e. DET w/ knob-adjustable strength would have been just as good as FRC). Maybe it would be better to have FRC combine DET and SPR rather than DET and NOFB.
- Spring had the steepest “learning curve” but I liked it a lot in the end. With practice, all could be usable. Maybe a basic system of visual cues; at least for training I needed to use my other hand to mark the physical location of the “start pitch” even though it was being played.
- I think DET condition made it more difficult to play in tempo as well. FRC seems to have the best of both worlds — offers precise controls (i.e. glissandi) as well as pitch guidance when force is applied.
- Cool stuff, takes a while to get hang of. DET harder to play glissando b/c easy to land on wrong notes. FRC much easier to guide.

- I am worried that I learned the melodies (and note positions) better as time passed during testing.
- I want one. It's like a Theremin but better. Have you seen the one (and only) old lady that is really good at the Theremin? I bet she was hot back in the day.

Appendix C

HSP Toolbox: An Overview

When we asked a colleague of ours why people do not make more haptic musical instruments, he replied that he thought they were “too hard to program and too expensive.” We decided to solve these perceived problems by introducing HSP, a simple platform for implementing haptic musical instruments. HSP obviates the need for employing low-level embedded control software because the haptic device is controlled directly from within the Pure Data (Pd) software running on a general purpose computer. Positions can be read from the haptic device, and forces can be written to the device using messages in Pd. Various additional objects have been created to facilitate rapid prototyping of useful haptic musical instruments in Pd. HSP operates under Linux, OS X, and Windows and supports the mass-produced Falcon haptic device from NovInt, which can currently be obtained for as little as US\$150. All of the above make HSP an especially excellent choice for pedagogical environments where multiple workstations are required and example programs should be complete yet simple.

C.1 Introduction

A *haptic musical instrument* consists of actuators that exert forces on the musician, sensors that measure the response of the musician, a program that determines what forces to exert on the musician, and a controller that executes the program and




Figure C.1: Musician interacting with a haptic musical instrument

interfaces with the sensors and actuators. Figure C.1 illustrates how the musician is included in the feedback loop, allowing the musician's gestures to be controlled. In most cases, the sensors and actuators are collocated to allow the control system to implement passive dynamical systems. Often the controller synthesizes sound signals in response to the way in which the haptic musical instrument is played.

C.2 Prior Work

Haptic musical instruments have been developed since the 1970's at the latest [41]. Much of the work has focused on the design of accurate haptic devices and physical modeling of virtual musical systems [42][45][63]. More recently, Sinclair [146] has controlled a SensAble Omni haptic device from Pd by employing elements from the CHAI 3D [51] and the Open Dynamics Engine (ODE) [121] libraries. ODE models the state of a virtual world, and CHAI 3D renders visual feedback and mediates the link between the virtual and the haptic worlds. However, the haptic feedback loop is not closed directly around messages passed in Pd; rather, Pd sends OSC messages to a haptics server that computes the forces to be applied to the haptic device. Currently CHAI 3D is fully supported only on Windows, but there is also a beta release for GNU/Linux.

We decided to instead develop our own platform because we wanted a truly cross-platform solution, and we wanted to process the haptic signals directly in Pd due



Figure C.2: Musician holding the Falcon grip

to its simplicity, instructional value, and the wide array of predefined objects for manipulating control messages in Pd. By employing this method for processing haptic control data, we cause ourselves to think using many of the computer music metaphors that motivated the development of Pd itself and its extended libraries.

C.3 Hardware

High resolution haptic devices have been becoming more accessible to musicians due to advances in gaming and minimally invasive surgery. We currently recommend that musical instrument designers use the NovInt Falcon, as shown in Figure C.2, which was originally designed for gaming. It currently costs about US\$150 and can deliver forces as large as 9 N. The position of the Falcon grip (see Figure C.2) is transmitted to a computer, and in response, the computer can command forces acting on the grip in 3D space. The Falcon communicates with the computer over USB at a servo rate of up to 1kHz.

Machulis' open-source driver for the Falcon operates under Linux, Windows, and OS X. There is no fully cross-platform open source driver for any other mass-produced haptic device capable of providing kinesthetic feedback. Machulis' flex external object *np_nifalcon* allows the Falcon to be controlled directly from Max/MSP or Pd




Figure C.3: Spring implementation in HSP

[105]. However, we provide a wrapper object called *Falcon* to hide some of the complexity of the underlying object, while still providing modularity for future updates of HSP allowing other haptic devices to be accessed.

It is important that the reader note how easy it is to install HSP since in contrast with many other approaches, nothing needs to be compiled and only one driver needs to be installed. The only required elements are the Pd patches we provide [14], the Falcon external object [105], and Pd extended.

C.4 1DOF Examples

C.4.1 Basic Elements

A motor induces a force F on the grip. The Falcon measures the movement of the grip in response as a position x . Hence, the most fundamental (i.e., memoryless) haptic algorithm for the Falcon implements a virtual spring with spring stiffness k .

$$F = -kx \quad (\text{C.1})$$

Figure C.3 shows how a one-dimensional spring can be implemented in Pd using the *expr* object, where the rest position of the spring is 0. Notice how simple the Pd patch is—only one object carries out any computations. This patch demonstrates how HSP’s simple philosophy of presenting the dynamics implementation directly leads to concise, easy-to-parse patches.

A related algorithm implements a wall at $x = 0$:

$$F_{wall} = -kx \cdot (x < 0) \quad (\text{C.2})$$

Another fundamental physical element is the damper. In order to simulate damping, the system needs to be able to estimate the velocity of the Falcon grip. The *viscous-friction-estimate* object is a simple estimator that is nevertheless fairly robust to noise and jitter. Rather than time-stamping each input position sample, we use the sample-and-hold Pd object *sig~* to upsample the position signal to the audio sampling rate, which is regular.¹ Next, we apply the filter $H(z)$:

$$H(z) = (1 - p) \frac{1 - z^{-1}}{1 - pz^{-1}}. \quad (\text{C.3})$$

For $p=0.0008$, $H(z)$ behaves like a differentiator from DC up until about 8Hz, at which point it rolls off the magnitude of the estimator signal. The roll off is necessary because velocity estimates are relatively noisy at higher frequencies. Finally, the velocity estimate signal is downsampled back to a control rate signal using the *snapshot~* object. The sampling period of the output signal is adjusted by changing the instantiation argument of *viscous-friction-estimate*.

C.4.2 Haptic Drum

The patch *drum.pd* simulates a wall, and whenever the musician presses into the wall, a drum sample is played back with an amplitude scaling proportional to the velocity. If the musician enables “ASSISTANCE”, then the instrument becomes unstable in a special manner enabling the musician to play especially fast drum rolls using only a single hand [27].

C.4.3 Detents and Force Profiles

Detents can help the musician orient himself or herself. Detents can be created even using 1DOF haptic interfaces. Figure C.4 illustrates how to implement a simple 1DOF detent. Near the center of the detent, the force profile looks like that of a spring, while the force goes to zero when the position y moves further from the detent’s center [166]. A simple potential energy argument implies that the force profile $F(y)$ is proportional

¹Note this means that audio must be enabled in Pd in order for the velocity estimator to function.




Figure C.4: Force profile $F(y)$ (above) and virtual terrain height profile $h(y)$ (below) for a simple detent.

to the derivative of the virtual terrain height $h(y)$ (see Figure C.4) implied by the haptic illusion, allowing various terrains and textures to be synthesized.

The GUI element bound to the *array* object in Pd provides for a convenient interface for exploring nonlinear force functions of position. Using springs in the X-axis and Z-axis, the Pd example patch *force-profile.pd* restricts the motion of the Falcon grip to the Y-axis, causing the Falcon to behave like a 1DOF haptic device. The force in the Y-axis is programmed to be a function of the Y-position. As illustrated in Figure C.5, the function is stored in a user-editable graphical array that can be modified with the mouse. At the center of the array, a horizontal slider (hslider) shows the Y-position of the Falcon (see Figure C.5, horizontal stripe in the middle). In this particular example, the user has drawn a force profile into the table resulting in two stable equilibrium points, toward which the Falcon grip is pushed by the motors. The stable equilibrium points have been marked in the image by way of small blue circles (see Figure C.5).

C.5 Multi-DOF Examples

C.5.1 Surface

A graphical array can also be used to model the height of a surface. The patch *surface.pd* uses a stiff spring in the X dimension to limit the Falcon grip to the Y-Z




Figure C.5: A force profile stored in a graphical *array* in Pd (stable positions circled in blue)

plane. Then whenever the musician presses into the surface, a spring force acts on the user in a direction approximately orthogonal to the slope of the surface above the Falcon grip. Friction may optionally be enabled that acts parallel to the slope of the surface and is proportional in magnitude to how far the user is pressing into the surface. Figure C.6 shows that a slightly unusual array configuration is chosen in this case to make it easier to obtain an estimate for the slope of the table as a function of position.

C.5.2 Models

A correctly implemented virtual model consisting of masses linked together by springs and dampers is guaranteed to not be unstable because there is no external source of energy. Similarly, the physical haptic device itself can be modeled approximately as a mass, so if it is connected by way of springs and dampers to a virtual model, the interconnected system will also not be unstable in theory. This implies that one meaningful and practical way of creating a haptic musical instrument is to use feedback control laws to link the haptic device to a virtual model by way of a spring [42] [2].

Consider the *bounce.pd* example in HSP, which employs the PMPD library for




Figure C.6: A surface modeled with a user-editable graphical array object in Pd

real-time modeling of virtual masses, springs, and dashpots vibrating at haptic rates [83]. The Falcon grip is connected to a virtual mass by way of a spring force, providing the impression that the Falcon grip has increased mass.

Acoustical virtual models are also of great interest so that sound can be synthesized directly from motion and/or force data according to the laws of physics. The open-source Synthesis ToolKit (STK), which can be accessed from Pd, provides many musical instruments such as stringed instruments, woodwinds, shakers, other percussion instruments, etc. that can be easily incorporated into a haptic musical instrument [54]. However, the *bowed-string.pd* example in HSP is implemented independently of STK. An acoustical model of a bowed string is constructed with delay lines and filters using the digital waveguide technique [149]. A collocated control junction is created in the center of the waveguide, from which point the virtual string is bowed according to a nonlinearity specified in a graphical array object. The musician can redraw the shape of the nonlinearity in order to alter the dynamics of the bow-string interaction.



Figure C.7: Falcon with shaker attached to grip

C.6 Performing Telemusic

Recently there has been interest in using networks to allow musicians at separate locations to play together in real-time. For instance, the SoundWIRE research group made it possible for musicians in Stanford, CA and Beijing, China to perform together in concert [39]. Audio is streamed over the network to allow performers to hear each other.

We suggest that with the help of haptic technology, a physical instrument may be teleoperated from a remote location. In other words, a musician may manipulate a haptic (master) device at one location, which is interfaced with a haptic (slave) device at a second location.

Figure C.2 shows the musician manipulating a Falcon at a machine running *telemusic1.pd*, while Figure C.7 shows a shaker attached to a Falcon at a machine running *telemusic2.pd*. The two patches communicate with one another over the network by way of the *netsend* and *netreceive* objects. The feedback control algorithm is designed so that a virtual spring links the two Falcons together. Since the spring is a bidirectional link, the musician can feel the mass of the shaker at the remote location, while the shaker at the remote location responds according to the actions of the performer.

C.7 Conclusions

HSP is an open-source platform that allows musicians to easily prototype haptic musical instruments. Since the haptic control signals are processed as messages within Pd, instrument designers are most likely to make use of the same metaphors that motivated the development of Pd itself and its extended libraries. For the convenience of the computer music community, HSP can be downloaded via the world wide web [14], and it is distributed under the GNU library General Public License.

We believe that HSP is simple to operate and that the examples included in the package demonstrate many of the ways in which haptics can be useful in musical instrument design. Accordingly, we have purchased six NovInt Falcons and created an instructional laboratory exercise around them for the CCRMA class entitled “Physical Interaction Design For Music” [15]. The example patches outlined above serve as excellent pedagogical material because they are simple and can be easily understood by students. In contrast with purpose-built platforms, HSP supports direct programming of the dynamics, so students can clearly comprehend the dynamics behind the physical interactions.

Appendix D

Nonlinear PID Control

D.1 Overview

Nonlinear proportional-integral-derivative (PID) control may be applied to change the dynamics of an acoustic musical instrument. While nonlinear network theorems from circuit theory may be applied by analogy to prove the uniqueness and stability of equilibria for more general nonlinear systems [172], we focus here on a special case because the stability proof is fairly simple. We choose the control input u to be the following:

$$u = P_D(\dot{x}, x)\dot{x} + P_P(x)x. \quad (\text{D.1})$$

Then substituting into (2.1), we obtain a new system incorporating a nonlinear damper and spring.

$$m\ddot{x} + (R + P_D(\dot{x}, x))\dot{x} + (k + P_P(x))x = 0 \quad (\text{D.2})$$

We collect the terms together, defining the functions $R(\cdot, \cdot)$ and $K(\cdot)$ to rewrite the nonlinear dynamics equation as follows:

$$m\ddot{x} + R(\dot{x}, x) + K(x) = 0. \quad (\text{D.3})$$




Figure D.1: Linear damping (left) and saturating damping (right)

D.2 Damper

The term $R(\dot{x}, x)$ in (D.3) describes damping. Although for the stability proof in Section D.4 it is not necessary that $R(\cdot, \cdot)$ be independent of x , many physical sources of damping are indeed independent of x . To gain some intuition into the nonlinear damping, we present some figures of damping where $R(\cdot, \cdot)$ is independent of x . Figure D.1 (left) shows the characteristic curve describing basic linear damping. For relatively small velocities, viscous damping tends to be linear [67]. Figure D.1 (right) shows the characteristic curve describing saturating damping, which bears similarity to Coulomb friction in brushed motors [56].

In this appendix, we introduce a more restrictive form of passivity that we call *Sastry-passivity*. Sastry-passive elements are indeed passive in the traditional energy sense from Section 2.3; however, Sastry-passive elements have additional properties allowing construction of the simple Lyapunov proof presented in Section D.4. Damping is *locally Sastry-passive* if $\dot{x}R(\dot{x}, x) \geq 0$ for all $x \in [-x_0, x_0]$ and $\dot{x} \in [-\dot{x}_0, \dot{x}_0]$ [172]. Damping is globally Sastry-passive if $x_0 \rightarrow \infty$ and $\dot{x}_0 \rightarrow \infty$. The damping force $R(\dot{x}, x)$ must act in the direction opposite the velocity \dot{x} . Hence, sources of Sastry-passive damping may only absorb and never source the power $\dot{x}R(\dot{x}, x)$.

Damping is *strictly locally Sastry-passive* if $\dot{x}R(\dot{x}, x) > 0$ for all $x \in [-x_0, x_0]$ and for all $\dot{x} \in [-\dot{x}_0, \dot{x}_0] - \{0\}$. Damping is strictly globally Sastry-passive if $x_0 \rightarrow \infty$ and $\dot{x} \rightarrow \infty$. In other words, the damping must be Sastry-passive and have




Figure D.2: Damping characteristic curve due to holding a stationary bow held against a string

no deadband. The linear and saturating nonlinear characteristics shown in Figure D.1 are both strictly locally Sastry-passive. Figure D.2 shows a characteristic curve simulating the effect of holding a stationary bow against a vibrating string [108]. In some models, the slope is chosen to be infinite at the origin, but here we choose the slope to be finite so that $R(\dot{x})$ is a function, allowing us to state that $R(\dot{x})$ is strictly Sastry-passive. Having finite slope at the origin is reported to incorporate effects due to the scattering of transverse waves into torsional waves at the bow [108].

D.3 Spring

The term $K(x)$ in (D.3) describes a nonlinear spring. A linear spring behaves according to $K(x) = kx$ for some constant k . Figure D.3 depicts a characteristic curve for a stiffening spring, which is quite common. At low frequencies, loudspeaker drivers exhibit stiffening spring behavior [99]. The stiffening spring or “tension modulation” effect of vibrating strings explains why the pitch of a vibrating string glides downward when plucked with a relatively large initial condition [81]. Chinese opera gongs are known for their pitch glide behavior too. Some of the gongs are designed to behave like stiffening springs, while others are designed to behave like softening springs (see p. 149 of [62]).

We similarly define Sastry-passivity of the spring analogously to Sastry-passivity




Figure D.3: Stiffening spring characteristic curve

of the damper [138][172]. $K(x)$ is *locally Sastry-passive* if $xK(x) \geq 0$ for $x \in [-x_0, x_0]$, and it is *strictly locally Sastry-passive* if $xK(x) > 0$ for $x \in [-x_0, x_0] - \{0\}$. The global definitions correspond to $x_0 \rightarrow \infty$. The stiffening spring shown in Figure D.3 is strictly locally Sastry-passive.

D.4 Lyapunov Stability Of Origin

We will show that the state $(x, \dot{x}) = (0, 0)$ is a locally asymptotically stable equilibrium point of the system described by (D.3) if the dashpot and spring are strictly locally Sastry-passive. Related proofs for other nonlinear networks are known from circuit theory [172].

First note that $(x, \dot{x}) = (0, 0)$ is indeed an equilibrium point because both sides of equation (D.3) are zero for $(x, \dot{x}) = (0, 0)$. We define the function $V(x, \dot{x})$ as follows [138][172]:

$$V(x, \dot{x}) = \int_0^x K(\sigma) d\sigma + \frac{1}{2} m \dot{x}^2 \quad (\text{D.4})$$

Because $K(x)$ is strictly locally Sastry-passive, $\int_0^x K(\sigma) d\sigma$ is locally positive definite about $x = 0$. $\frac{1}{2} \dot{x}^2$ is positive definite in \dot{x} . It follows that $V(x, \dot{x})$ is locally positive definite [138]. Taking the time derivative of (D.4), we obtain \dot{V} :

$$\dot{V}(x, \dot{x}) = -R(\dot{x}, x) \dot{x}. \quad (\text{D.5})$$

Since the dashpot is strictly locally Sastry-passive, $\dot{V}(x, \dot{x}) \leq 0$ in the neighborhood of the origin. The only trajectory in the neighborhood of the origin where $\dot{V}(x(t), \dot{x}(t)) = (0, 0)$ for all t is the trivial trajectory $(x(t), \dot{x}(t)) = (0, 0)$. By LaSalle's Theorem, $(x, \dot{x}) = (0, 0)$ is a locally asymptotically stable equilibrium point [138].

The proof above shows that if we wanted to control a musical mass-spring-damper system, if we applied a nonlinear damping control force and nonlinear spring force, then the system state will decay asymptotically to $(x(t), \dot{x}(t)) = (0, 0)$ under the following conditions:

- The net damping force is strictly locally Sastry-passive.
- The net spring force is strictly locally Sastry-passive.
- The initial state is chosen sufficiently close to the origin $(x(t), \dot{x}(t)) = (0, 0)$.

D.5 Implications Of Passivity

In a practical musical application, it is perhaps more convenient to appeal directly to passivity theory rather than Lyapunov theory because then we do not need to find specific Lyapunov functions. In particular, let us assume that the acoustic musical instrument being controlled is dissipative. It can have as many modes of vibration as desired. Then as long as the controller is passive, which is less restrictive than the controller being Sastry-passive, it cannot add any energy to the instrument. This implies that the net energy in the controlled system cannot increase faster than the musician adds it, so the dynamic behavior is not unstable.

In the laboratory, we have implemented nonlinear damper and nonlinear spring controllers. Various effects can be obtained while restricting the damper and spring to be passive. For example, the nonlinear spring controller may be applied to create pitch glide effects as with the gongs described in Section D.3. The passive nonlinear damper controller causes string vibrations to decay while causing buzzing or rubbing sounds (e.g. see Figure D.1, right).

Appendix E

Additional Musical Instruments

E.1 Tangible Virtual Vibrating String

E.1.1 A Physically Motivated Virtual Musical Instrument Interface

We introduce physically motivated interfaces for playing virtual musical instruments, and we suggest that they lie somewhere in between commonplace interfaces and haptic interfaces in terms of their complexity. Next, we review guitar-like interfaces, and we design an interface to a virtual string. The excitation signal and pitch are sensed separately using two independent string segments. These parameters control a two-axis digital waveguide virtual string, which models vibrations in the horizontal and vertical transverse axes as well as the coupling between them. Finally, we consider the advantages of using a multi-axis pickup for measuring the excitation signal.

E.1.2 Introduction

Physical Models

Virtual models of acoustic musical instruments have been available to the music community for decades [41] [149] [96]. The models are useful for studying the physical behavior of acoustic musical instruments, and they can also synthesize sound output.

Given an appropriate interface, many of the models can be played in real-time by performers.

Commonplace Interfaces Often it is most convenient and simplest to control a physical model with a commonplace interface, such as a computer keyboard, musical keyboard, or mouse. This approach is most palatable if the interface matches the physical model. For instance, playing a virtual piano with a musical keyboard interface is physically intuitive, so it is easy for a pianist to transfer real-life skills to the virtual domain. However, many performers play traditional acoustic instruments lacking commonplace interface counterparts, so skill transfer to the virtual domain is not as immediate [110].

Haptic Interfaces Haptic interfaces lie at the opposite end of the complexity spectrum. They apply force feedback to the performer, so that he or she feels and interacts with the vibrations of the virtual instrument as if the virtual instrument were real. In this sense, haptic interfaces can be seen as the ideal interface for interacting with virtual instruments. For instance, a carefully designed haptic bowed-string should promote better skill transfer to the virtual domain because it exerts forces on the instrument interface causing it to behave as if it were a real bow bowing a string.

When Luciani et al. implemented their haptic bowed string, they found that users strongly preferred that haptic feedback be rendered at the audio rate of 44kHz rather than at the usual 3kHz. Users made comments regarding the “strong presence of the string in the hand,” “the string in the fingers,” and “the string is really here” [63]. Related kinds of instruments, such as actively controlled acoustic musical instruments are essentially the same as haptic musical instruments except that the whole acoustical medium becomes the interface [20]. Haptic technologies are becoming increasingly available to the music community, but they are currently still complex enough that it is worth considering alternatives.

Physically Motivated Interfaces In this paper, we investigate the middle ground in between commonplace interfaces and haptic interfaces for controlling physical models. We term such interfaces *physically motivated interfaces*. Rather than applying haptic feedback, we attempt to otherwise preserve the physical interaction between the performer and the virtual instrument as much as possible. Such interfaces are

similar to Wanderley’s categorization of instrument-like controllers with one important exception [110]: we state that the input quantities should be sensed so accurately that an audio-rate feedback loop could be closed around the sensor if the interface were equipped with an actuator. It follows that ideally all quantities applied to the physical model should:

- correspond to the correct quantity for controlling the physical model (e.g. displacement, velocity, acceleration, etc.)
- be linear and low-noise
- be delayed and filtered as little as possible
- be sampled at the audio sampling rate

These requirements were difficult to meet in the past due to limitations in computational power, A/D conversion, and sensing; however, today we may achieve or approximate them as we see fit. To succeed in our endeavor, we need to carefully apply knowledge from acoustics, mechanical engineering, and electrical engineering to the field of human computer interaction. In the following, we develop a physically motivated interface for a virtual vibrating guitar string.

E.1.3 Prior Guitar-Like Interfaces

A number of musical instrument interfaces suggest the metaphor of a guitar. While they have followed different design goals, we should at least consider how they estimate the desired pitch. For instance, the virtual air guitar uses the distance between the hands to control the pitch. Different versions of the virtual air guitar make use of magnetic motion capture, camera tracking, and acoustic delay estimation systems for this measurement [95].

The makers of the GXtar prefer to place an force-sensing resistor strip placed beneath a real string to measure both the position and pressure [97]. The Ztar [150] and the Yamaha EZ-GE [174] detect pitch with a matrix of sensors in the neck. One sensing element is used for each fret and string. The SynthAxe sports normal strings

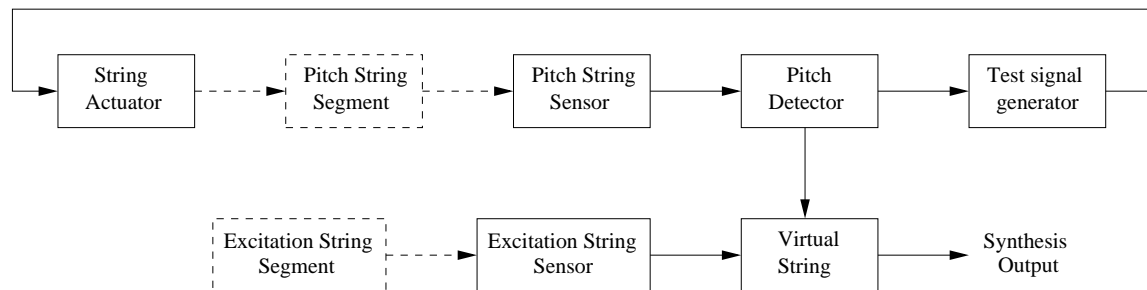


Figure E.1: Tangible virtual string signal flow diagram

placed above matrix-type sensors [3]. The SynthAxe has additional sensors to detect string bending. A current flows down the each string, and small electric coils placed near the string measure the lateral string displacement.

Roland provides a six-channel electromagnetic pickup for electric guitar and accompanying DSP [130]. In the MIDI mode of operation, a DSP estimates when new notes are played, with what velocity, and with what pitch.¹ One drawback of this approach in general is that most detectors have noticeable delay for lower pitches because they wait at least one period. It is also difficult to construct a perfectly reliable pitch detector of this type. Consequently, performers must learn to play carefully to avoid confusing the pitch detector.

E.1.4 Tangible Guitar String

Separate Excitation Sensing and Pitch Detection

To ensure that the interface is physically motivated, we follow the guidelines outlined in Section E.1.2. We sense the relevant portions of performer’s gestures with as much precision as possible to preserve the guitar-like physical interaction between the performer and the virtual instrument. We would also like to preserve the physical presence of a string. To these ends, an independent string segment associated with each hand separates the problems of estimating the desired pitch and measuring the

¹In another mode of operation, the Roland system synthesizes audio more directly from the sensed signals using “Composite Object Sound Modeling”. Here the pitch detector is not needed explicitly, so tracking is much improved. Since the model is not entirely virtual and its details are trade secret, we do not consider it further here.

plucking excitation signal [3].

Signal Flow

The upper half of Figure E.1 shows the signal flow for the pitch string segment; the purely acoustic components and paths are drawn in dashed lines. We detect the desired pitch of the string acoustically to help avoid incorrectly capturing higher order effects such as string bending, slightly misplaced frets, etc. We actuate the string and measure its response. Since we know how the string is being actuated, we should be able to more accurately estimate the length of time it takes for a pulse to leave the actuator, reflect off of a fret, and arrive back at the sensor (see Figure E.2, top).

Any picking, plucking, scraping, or bowing excitation is sensed via the excitation string segment and fed directly to the virtual string.² The lower half of Figure E.1 shows the signal flow for the excitation string segment. One end of the excitation string segment should be damped passively to prevent physical resonances from interfering with resonances in the virtual model (see the damping material in Figure E.2, bottom).

Two-Axis Digital Waveguide Virtual String

We model the virtual string using a simple two-axis model that takes into account the vertical and horizontal transverse modes of vibration. The i th axis is modeled using a delay line of length N_i samples and lowpass filter LPF_i , which is a 3-tap linear-phase FIR filter causing the higher partials to decay faster (see Figure E.3). This portion is the basic digital waveguide model used for elementary teaching purposes at CCRMA. For additional realism, the excitation signals can be comb filtered with the notch frequencies chosen as a function of the excitation's distance from the bridge [149].

²There are surprisingly few examples in the literature where some filtered form of an excitation signal measured at the audio sampling rate appears at the output. One example is the digital flute, which allows the excitation signal as measured by a microphone to be mixed with the sound synthesis output; however, in contrast with the current work, sound was not synthesized with a physical model [176].




Figure E.2: Two string segment approach

While the nut is assumed to be rigid, the bridge is in general not quite rigid, hence it couples the axes together at this point. The coupling implemented in Figure E.3 is actually more appropriate for modeling the coupling of the vertical axes of two neighboring piano strings, but it still results in qualitatively correct behavior. For instance, by choosing $N_1 \approx N_2$ such that $N_1 \neq N_2$ and $g \approx 0.1$, one obtains behavior where the energy slowly rotates back and forth between the axes of vibration [149] [102].

Prototype

A prototype of the tangible guitar string interface is shown in Figure E.5. For convenience given the default hardware on a Fender Stratocaster, the two string segments are spaced horizontally instead of vertically in relation to one another. In a six-stringed embodiment, each pair of string segments would instead be placed axially-aligned with each other. In the prototype, each string's vibration is sensed using a Graphtech saddle piezoelectric pickup, as shown in Figure E.4 [70].

The magnetic actuator can be obtained by ordering the Sustainiac [84]. It conveniently replaces any one of the pickups; however, other kinds of actuators can be used

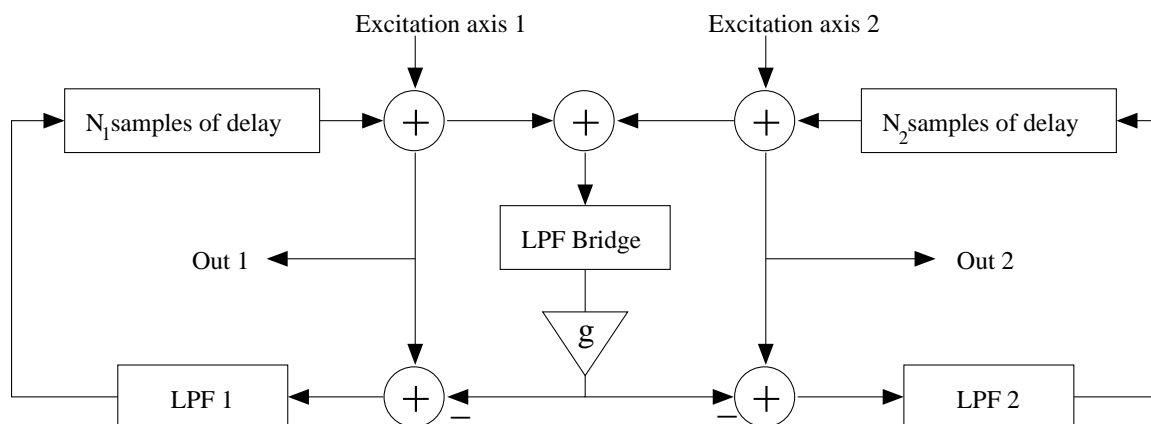


Figure E.3: Two axis digital waveguide string model



Figure E.4: Graphtech piezoelectric pickup

instead. To prevent the actuator from affecting the excitation string segment, we choose the excitation string segment to be a regular, non-ferrous solid electrical wire. The excitation string is passively damped using felt, wrapped in such a manner to approximate gradually increasing the string's wave impedance to infinity, eliminating reflections as much as possible. In other words, the strip of felt is wrapped more and more tightly approaching the nut (see Figure E.5). In some cases, it is better to damp the string less effectively. The resulting less damped reflections from the felt material cause a comb-filtering effect, which changes the timbre of the instrument as a function of excitation position as with a normal vibrating string. Note that this desirable attribute comes for free since the interface is physically motivated—the comb filter can be implemented either mechanically on the interface or virtually in the instrument model.

Multi-Axis Pickups To most accurately excite the physical model, we should ideally measure the excitation in both the horizontal and vertical transverse axes of the

string. This can be done using optical [80], electromagnetic, or piezoelectric sensors [66]. We have verified informally that the tangible virtual string sounds more realistic given the two-dimensional excitation.

E.1.5 Website

We have authored a website providing sound examples of the instrument in a few different configurations.³ For example, the website includes comparisons of model output given randomly synthesized excitations, single axis measured excitations, and two-axis measured excitations. It also includes model output given various excitation sources such as plucking, picking, bowing, and scraping. To enable others to excite their physical models with quality excitation signals, we provide the corresponding non-resonant excitation signals themselves. Finally, an example melody played on the tangible virtual string demonstrates the viability of physically motivated instrument design.

E.1.6 Future Work

The behavior of the interface could be further refined with force-feedback. For example, the excitation string segment could be made into a haptic device by adding an actuator. Then the piece of physical string could be joined to a portion of the waveguide using teleoperator techniques [79]. It would be essential that the string segment would have as little mass as possible to avoid loading down the virtual waveguide at the point of connection. We would also like to eventually construct a six-string version to promote the maximum transfer of guitarists' skills from real guitars to virtual guitars.

We currently excite the pitch string segment with filtered noise, and we estimate the pitch using `fiddle~` in Pure Data [132]. While this is sufficient for a prototype, improved versions would benefit from superior pitch tracking. `fiddle~` neither makes use of knowledge of the excitation signal, nor does it take advantage of any optimizations given physical knowledge about vibrating strings. For example, we would like

³<http://ccrma.stanford.edu/~eberdahl/Projects/TS>

to consider a time-domain delay-based method, similar to the method explained by Karjalainen [95].

E.1.7 Conclusion

We have presented a physically motivated interface for controlling a virtual digital waveguide string. The excitation and pitch are sensed separately using two independent string segments. In contrast with prior interfaces, the excitation to the physical model is measured according to the principles of *physically motivated* interfaces. In particular, we measure the excitation signals with high quality, linear, and low noise sensors at the audio sampling rate. We hope that interfaces such as this one will continue to promote skill transfer from traditional acoustic musical instruments to the virtual domain.

E.2 Electromagnetically Prepared Piano

In an attempt to create alternative methods of both playing and studying vibrating strings, we have constructed a software-driven instrument for use in the electromagnetic excitation of an acoustic grand piano’s strings.⁴ Applications to musical composition and expression, piano string characteristic identification, arbitrarily long sustained tones, and digital waveguide model calibration are presented, along with ideas for future experimentation and creation of new music and sound.

Composer Per Bloland initially motivated this research by requesting a device capable of vibrating piano strings using electromagnetic waves. Commissioned by SEAMUS/ASCAP 2005, Bloland’s acoustic piano composition entitled *Elsewhere is a Negative Mirror, Part I* incorporates an electromagnetic preparation of the piano.

Although the final prototype turned out to be much different, in some ways the early conceptions of this instrument gave a nod to the popular “Ebow” [82] string sustainer long used by electric guitarists and others. Perhaps the most significant deviation from this design is that we do not rely on mechanical-electrical feedback

⁴This section is taken almost verbatim from a paper written by Edgar Berdahl, Steven Backer, and Julius O. Smith III [16].

in the same sense; instead, we transmit sounds into the piano, which then emanate acoustically from the piano without any internal electrical feedback. Possibly the largest similarity is that the sound of the piano's strings and body can be made to endure for long periods of time, a quality found in many extended bowing techniques.

Electromagnetism has long been used in musical instruments and related equipment. Of course, it can be found in virtually any computer music environment, whether introduced intentionally or inherently. Conventional speakers, microphones, hard disk drives, and other means of electronically manipulating and producing sound rely on electricity and magnetism. In addition, many composers and scientists have explicitly exploited electromagnetism. Maggi Payne's *Holding Pattern*, Stephen Scott's *Resonant Resources*, and John Cage's *Postcard from Heaven* are just a few that come to mind. Yet another relevant reference to magnetic vibration of strings is the work done by Weinreich and Caussé [168] involving bowed string motion simulation using hybrid mechanical-electrical systems. Our excitation device is capable of receiving an arbitrary waveform from any common computer sound card or electrical signal source, and in turn relaying this information to a piano string, without physically making contact with the string itself. Widely available software such as Pure Data or Cycling74's Max/MSP is ideal for sending a variety of signals to the device, from pure sine tones, rich orchestral samples, and voice clips, to simply white noise. However, due to the flexibility of the device, one need not be constrained to any specific software or type of sound.

From a hardware prototype point of view, the system is quite simple; the three major building blocks consist of an audio power amplifier, an electromagnet, and two permanent bar magnets positioned orthogonally to the string lying in a plane a short distance above the string (see Figure E.6). Combined in the right manner and within the physical constraints presented by the piano, arbitrary string motion can be created to a degree of accuracy high enough for both scientific measurements and subjective aural observations (see Section F.2.2).

Sustained resonances can be produced from strings over almost the entire range of the piano. Also, individual partials of each string can be isolated and evoked independently of the fundamental string vibration frequency. In some cases one can

even reproduce a continuum of sound over frequency from a single string. Some sound samples are posted on a website,⁵ though they are by no means representative of the full capabilities of the instrument.

In the block diagram shown below (see Figure E.7), the basic structure of the device is depicted for one channel.⁶ Beginning as a concept in computer software, a musical signal finds its way from the output of a sound card (DAC) to the input of an audio amplifier via standard audio cables and jacks. Instead of connecting a conventional speaker to the amplifier's output, a similar load was presented in the form of an electromagnetic transducer with a real impedance of 8Ω . The portion of the device that rests within the piano was mounted on a piece of oak laid atop the large crossbars that connect to the piano frame (see Figure E.8). The electromagnets themselves were originally intended for holding applications.

⁵<http://ccrma.stanford.edu/~sbacker/empp/>

⁶A total of 12 channels / notes were implemented.



Figure E.5: Tangible Guitar String Interface



Figure E.6: Close up of the motors




Figure E.7: System block diagram for one channel



Figure E.8: Electromagnetically prepared piano prototype

Appendix F

Actuator and Sensor Design

The treatment here covers only sensors and actuators that we have used in the laboratory, and for this reason, it is by no means complete. Nevertheless, we believe that readers wishing to construct their own sensors for applications in acoustics, and particularly in musical acoustics, will find the information here helpful in practice, especially as there are so few references in the literature.

F.1 Harmonic Distortion

Harmonic distortion can be a problem for both actuators and sensors. We define the total harmonic distortion (THD) in this section so that we can refer to it consistently in other sections. Consider the single-input single-output (SISO) system shown in Figure F.1. Let the input $u = A \cos \omega t$. Most deterministic systems¹ will produce output of the form

$$y = \sum_{n=0,\dots,\infty} B_n \cos(n\omega t + \phi_n). \quad (\text{F.1})$$

We define the total harmonic distortion k to be

¹Exceptions include chaotic systems and systems exhibiting subharmonic behavior.



Figure F.1: SISO system

$$k = \sqrt{\frac{\sum_{n=2,\dots,\infty} \frac{1}{2} B_n^2}{\sum_{n=1,\dots,\infty} \frac{1}{2} B_n^2}} \quad (\text{F.2})$$

and the partial harmonic distortion components k_n to be

$$k_n = \left| \frac{B_n}{B_1} \right| \quad (\text{F.3})$$

for $n = 2, \dots, \infty$ [157].

F.2 Actuators

In actuator design, there is typically a trade-off between nonlinearity and maximum displacement [99]. Since the power exerted by a transducer is often related to the displacement, we may also state that there is a typically a trade-off between nonlinearity and power. We consider two types of string actuators in the following sections.

F.2.1 Lorentz' Law Type

The Lorentz' law type of string actuator behaves simply. Figure F.2 shows a diagram of the geometry. A conductive vibrating string is placed between two permanent magnets. For simplicity, we assume that the magnetic field \mathbf{B} is completely uniform in between the magnets, and that it never flows back to complete the magnetic circuit (see Figure F.2). This is of course impossible, but in reality the field flowing from the north pole of the upper magnet to the south pole of the lower magnet is much less focused. Consequently, it can be neglected to first order.²

²Some small portion of this reverse field does indeed intersect the string outside of the main actuator window, and will be responsible for a slightly high-pass actuation transfer function.




Figure F.2: Lorentz' law type actuator

Lorentz' law states that a wire with a current I flowing along it in the presence of a magnetic field \mathbf{B} of length L_A will experience the force $\mathbf{F}_{\text{string}}$, as given in (F.4). According to the geometry shown in Figure F.2, this force points out of the page. An important feature of this actuator is that for large magnets, \mathbf{B} is nearly uniform over the volume intersected by the vibrating string, so the actuator behaves especially linearly. The force on a loudspeaker cone is also generated according to Lorentz' law [99] [65].

$$\mathbf{F}_{\text{string}} = L_A \mathbf{I} \times \mathbf{B} \quad (\text{F.4})$$

A piece of wire of only about 1m in length has a relatively low resistance. In order to connect it to the output of a typical audio power amplifier, it must be placed in series with power resistors to avoid overloading the amplifier's output. Most of the power is dissipated by the power resistors R rather than in doing work on the string, making this actuation scheme inefficient and wasteful in terms of power. However, since the circuit connected to the output of the amplifier consists of only one loop, whose impedance is dominated by power resistors R , the load is seen as purely resistive. This property is useful because the almost all audio power amplifiers induce a voltage V across their output terminals. However, the force on the string is a function of the current I flowing along the string. Due to Ohm's law ($V = IR$), we have that $V \propto I \propto F_{\text{string}}$, meaning that we do not need to worry about the power amplifier and actuator being responsible for additional poles or zeros in the control




Figure F.3: Lorentz' law type actuator (viewed from above)

loop.³

Consider that if the string is ferrous, it will be attracted to the magnets to a varying degree depending on its position. This nonlinear spring force causes a slight detuning of the string's natural harmonic series. This effect is known colloquially in the electric guitar community as “stratitis” after the Fender Stratocaster electric guitar. Stratitis occurs when guitarists place their pickup's magnetic pole pieces too closely to the string. We avoid stratitis in the laboratory by using a nonferrous string, made out of a standard piece of 18AWG solid wire. This type of vibrating wire nevertheless sounds very similar to a standard musical vibrating string.

The realization of the Lorentz' force actuator in the laboratory is shown in Figures F.3 and F.4.

F.2.2 Reluctance Force Type

A ferrous vibrating string can be actuated by a reluctance force actuator. This is the kind of actuator found in both the E-Bow [82] and Sustainiac [84] electric guitar string vibration sustainers. A hand-wound reluctance force actuator is shown in Figure F.5. It consists of a coil wound around a magnetically permeable core, for instance iron. In Figure F.5, the core is a pole piece removed from a Fender electric guitar pickup.

³Of course, all equalizing filters in the power amplifier need to be disabled. Most audio power amplifiers also place zeros at DC, but luckily these do not usually interfere with actively controlling the string—they affect the closed-loop transfer function only far below the first resonance of the musical system.




Figure F.4: Lorentz' law type actuator (oblique angle)




Figure F.5: Hand-wound reluctance force actuator

A neodymium permanent magnet is placed at the end of the core to set up a strong DC field. The DC field magnetizes the ferrous string in the neighborhood of the pole piece, so that small currents flowing through the electromagnet can push and pull on the string [144]. A heat sink helps dissipate heat generated in the coil. The pole piece, magnet, and heat sink are held together using thermal epoxy.

The goal the following analysis is to more rigorously determine the relationships between the force f_X on the element in motion, the current $i(t)$ flowing through the coil, and the air gap x , which is the vertical distance between the string and the transducer.⁴ In general, f_X and $i(t)$ will be nonlinearly related, as depicted in Figure F.6 (right). f_X will be approximately proportional to $i^2(t)$, except for fields so

⁴The remainder of Section F.2.2 is taken almost verbatim from a paper written by Edgar Berdahl, Steven Backer, and Julius O. Smith III [16].

large that the string saturates magnetically, in which case f_X becomes nearly linearly related to i (circle in Figure F.6). By placing the permanent magnet such that the field is focused on the string, we can make the transducer operate in an approximately linear region.

The force is commonly also a hyperbolic function of x (Figure F.6, left). This is undesirable. We have made measurements showing that when string displacements are relatively small, the system is roughly time-invariant [16]. Unfortunately our analysis is not so simple as in [144] because the path length that the magnetic flux flows along through the string is not constant. However, we will show that the analysis holds for an arbitrarily small string element.




Figure F.6: Relationships between f_X , i , and x

We can simplify the analysis by using the superposition principle to describe the flux density $\mathbf{B}(t)$ at time t at a particular point on the string. That is, similarly to the linearized analysis of analog circuits containing transistors, the quantity $\mathbf{B}(t)$ can be split into a large component $\mathbf{B}_L(t)$ and a small component $\mathbf{B}_S(t)$.

$$\mathbf{B}(t) = \mathbf{B}_L(t) + \mathbf{B}_S(t) \quad (\text{F.5})$$

$\mathbf{B}_L(t)$ is due to a magnetic “biasing” of the transducer-string system so that it operates nearly linearly. $\mathbf{B}_L(t)$ is due to the permanent magnet. $\mathbf{B}_S(t)$ is proportional to $i(t)$ [91]. $\mathbf{B}_L(t)$ magnetizes the piano string so that it can be more easily acted upon by the coil. However, because the magnetic field due to the permanent magnet

is much stronger than that due to the coil, we will neglect the coil's contribution here.

$$\mathbf{B}_L(t) = \mathbf{B}_{\text{pm}} + \mathbf{B}_{\text{coil}}(t) \approx \mathbf{B}_{\text{pm}} \quad (\text{F.6})$$

The qualitative plot in Figure F.7 shows the shape of the magnetic field lines \mathbf{B}_L as approximated by a two-dimensional electromagnetic field simulator. However, the geometry is for a different actuator with a wider core and multiple permanent magnets as compared to the geometry in Figure F.5, but the fields should be similar because each geometry contains a core that is magnetized along the same axis. The field lines tend to flow along the string rather than near it because the magnetic permeability of the string μ is much higher than that of free space μ_0 [91].




Figure F.7: Expected magnetic field lines for a different geometry (wider core, multiple permanent magnets)

An arbitrarily small string element near the transducer can be considered to be roughly constantly magnetized since the core of the core is much larger in diameter than the string, since the variations in x are small relative to x , and since \mathbf{B}_{pm} is approximately constant. We shall thus consider the magnetic moment \mathbf{m} of the string element to be roughly constant.⁵ Now that the portion of the string near the permanent magnets is magnetized, any change in the magnetic field will either push or pull on the string depending on the direction of the change. The current

⁵For larger displacements, \mathbf{m} will vary more in practice. For this reason, reluctance actuators are more linear when applied to piano strings in comparison with guitar strings, which generally vibrate over larger displacements.

$i(t)$ flowing through the coil creates small time-varying changes in the magnetic-flux density $\mathbf{B}_{\text{coil}}(t)$ at the string element we are considering.

$$\Delta \mathbf{B}(t) \triangleq \mathbf{B}_S(t) = \mathbf{B}_{\text{coil}}(t) \quad (\text{F.7})$$

A magnetized element in the presence of an external magnetic field will experience a force. In our case, the magnetized element is the string element, and the external magnetic field is $\mathbf{B}_{\text{coil}}(t)$. Let $U(t)$ be the potential energy of the string element at time t that will contribute to the string's vibration.

$$U(t) = -\mathbf{m} \cdot \mathbf{B}_{\text{coil}}(t) = -\cos(\theta)|\mathbf{m}||\mathbf{B}_{\text{coil}}(t)| \quad (\text{F.8})$$

The force in the x -direction can be determined by taking the derivative with respect to the x -direction.

$$f_X = -\frac{\partial U(t)}{\partial x} \quad (\text{F.9})$$

Since the flux density $\mathbf{B}_{\text{coil}}(t)$ is proportional to the current and roughly inversely proportional to a positive power M of the air gap x , we arrive at the proportionality relationships that were previously depicted in Figure F.6:

$$f_X \propto x^{-M-1} \quad (\text{F.10})$$

and

$$f_X \propto i(t). \quad (\text{F.11})$$

Reluctance force actuators are more practical than Lorentz' law actuators because there is no current flowing along the string, hence the musician does not need to be concerned about being shocked. Installing such actuators is also minimally invasive. In fact, the E-Bow makes this aspect part of the control interface. The musician holds the E-Bow in the neighborhood of the string that the musician wants to start vibrating [82].

One disadvantage of reluctance force actuators is that the relatively large number

of windings causes the actuator to have a significant inductance L , which may for example be on the order of a few millihenries. Thus, although in the linearized model the current $i(t)$ is proportional to the force on the string f_X , an audio power amplifier will induce a voltage V across the terminals of the actuator. If we label the DC resistance of the coil R , then we can determine the transfer function between the commanded voltage $V(s)$ and the current $I(s)$. The inductance is responsible for an additional pole in the feedback loop as given in (F.12), which must be considered when developing controllers. Besides moving this pole to a higher frequency with a lead, we can consider applying an alternate kind of power amplifier, known as a constant-current amplifier.

$$\frac{V(s)}{I(s)} = \frac{1}{L} \cdot \frac{1}{s + \frac{R}{L}} \quad (\text{F.12})$$

F.3 Sensors

There are many different kinds of sensors, so in this section we focus on some that are useful for sensing vibrating string displacements. For a given type of sensor, there is a trade-off between nonlinearity and noise.

F.3.1 Optical Displacement Sensors

PHOTODIODES

Optical displacement sensors can directly measure the displacement of a vibrating string at a point along the string. Since they work all the way down to DC, they are relatively easy to test and calibrate. For example, consider the configuration shown in Figure F.8, which is recommended by Weinreich and Causse [168]. An infrared photodiode is masked with an infrared-reflecting material, such as aluminum, so that a triangular sensitive area is exposed. The photodiode is then placed as close to a vibrating string as is reasonable.

An idealized representation is shown in Figure F.9, where an infrared light source (not shown) illuminates the sensitive surface from behind the string. The string casts a shadow on the triangular surface, hence the amount of light detected is affinely




Figure F.8: Photodiode type optical sensor in the laboratory (LED not shown)




Figure F.9: Photodiode type optical displacement sensor

related to the motion of the string in the plane parallel to the sensitive surface.

We have constructed this type of sensor using two different kinds of photodiodes. The S2387-1010R silicon photodiode from Hamamatsu has a large sensitive surface of 1cm by 1cm (see Figure F.8). The large size is desirable because it allows large string displacements—for example, it can be placed at any position along a vibrating guitar string. However, we have had little success at implementing a linear sensor with this photodiode. The smallest THD we were able to obtain for vibrations of about 3mm in peak-to-peak amplitude was -35dB. The nonlinear performance is likely partly due to the slightly nonuniform sensitivity of the photodiode near the edges of the sensitive

surface.⁶ We could have mitigated this problem by also masking the edges of the photodiode. However, we believe that the poor performance was due to difficulties in uniformly illuminating the relatively large sensitive portion of the photodiode surface with a sufficient amount of infrared light.

In practice, we used the PD638B silicon photodiode (not shown) from Everlight. While the diode's sensitive surface measures only about 3mm x 4mm, it is easy to obtain, and it includes an infrared filter, greatly reducing the effects of ambient visible light in the room on the measurement. To convert the photodiode current into a voltage, we used the generic transimpedance amplifier circuit described by Graeme and many others [69]. We chose the resistor and capacitor in the negative feedback loop of the circuit to be $R = 10k\Omega$ and $C = 0.7nF$, implying a theoretical 3dB bandwidth of $f_{3dB,theoretical} = 21.5kHz$.⁷

LIGHT SOURCES

Light emitted from a flat surface element is subject to Lambert's cosine law, which is also sometimes known as the cosine law of illumination. The light intensity emitted is proportional to the cosine of the angle formed with the surface element normal [32].⁸ In many situations, light is emitted from a surface, which explains one instance of the sensing trade-off between non-linearity and noise. For a given light source with non-uniform light intensity radiation pattern, moving the light source further away from the sensing element makes the illumination of the sensing surface more uniform. However, as the amount of light cast upon the surface decreases, the sensor noise remains constant. Thus, moving the light source further away from the

⁶Christine Nishiyama, an applications engineer for Hamamatsu, stated in personal communications that "the uniformity is within 1% at 80% of the entire active area. At the outer edge of the active area is where the linearity gets worse."

⁷We made a measurement to verify the actual achieved sensor bandwidth $f_{3dB,measured}$. We drove an infrared LED directly with a sinusoid from a signal generator, and we measured the 3dB bandwidth of the operational amplifier's output voltage to be about $f_{3dB,measured} \approx 8.5kHz$. Replacing our TL072A operational amplifier with other operational amplifiers having larger gain-bandwidth products did not change $f_{3dB,measured}$. More complicated modeling of the sensing circuit would undoubtedly explain why the actual bandwidth was smaller, and much more complicated models are available [69], but we considered $f_{3dB,measured} \approx 8.5kHz$ to be large enough for our application.

⁸Xenon arc lamps are not subject to the same limitation because the light source is due to an arc; however, xenon arc lamps with power supplies generally cost more than \$2000 each because they target purely scientific applications.




Figure F.10: Halogen lamp

sensing element decreases the signal-to-noise ratio.

The most practical source of infrared light is an infrared LED. Such LEDs are small, inexpensive, and allow easy placement of multiple sensors along the string. However, they do not have a uniform light radiation pattern. In the laboratory, we use the high intensity infrared LED IR204/H16/L10 from Everlight. The character of its radiation pattern conforms qualitatively to Lambert's cosine law, although the LED surface is not flat [32]. For this particular LED, the radiant intensity falls off by only 6% by 10 degrees from the normal [160]. Although the maximum power dissipation of the LED is only 150mW, we found that placing the LED about 2cm from the PD638B silicon photodiode made for a fairly good displacement sensor. The signal-to-noise ratio was good enough for most control experiments. When carefully aligned, the THD was about -40dB for string displacements of 2mm in peak-to-peak magnitude.

Halogen lamps are worth considering because they produce so much output power. For instance, consider the 20W halogen lamp shown in Figure F.10. However, the bright element consists of a coil, and the coil loops cause shadowing effects producing a very nonuniform radiation pattern, which is illustrated in Figure F.11 by holding a white sheet of paper behind the lamp.

Nevertheless, Figure F.12 reveals that by orienting a halogen lamp so that the line through the center of the coil is a few degrees off axis with the flat surface to be illuminated, an approximately uniform region of illumination is created, as evidenced inside the red loop in Figure F.12. We found that using this light source in conjunction with either of the photodiodes previously discussed greatly improved the



Figure F.11: Nonuniform radiation pattern of a coiled halogen lamp

noise performance. In fact, we would prefer using the halogen lamp in conjunction with the smaller PD638B photodiode because the signal-to-noise ratio (and inductive coupling from the actuator) was improved by about 20dB. However, our original circuit design was not motivated by heat considerations, and the large power level of the light source caused us to reconsider the practicality of using an infrared LED as the light source. That is, for our experiments in Chapters 3 and 4, we used the PD638B photodiode in conjunction with the IR204/H16/L10 LED.

PHOTOTRANSISTORS

The sensing circuit for phototransistor type measurements is simpler. Rather than requiring an additional operational amplifier as in the case of photodiode sensors, the transistor-like properties of the sensing element may be used directly; hence, the amplifier circuit consists only a voltage source, a few resistors, and the phototransistor [80]. Because the phototransistor produces such a large voltage output, the measurement is likely less susceptible to inductive coupling effects (see Section F.4).

The bandwidth of phototransistors is typically more limited in comparison with photodiode sensors. For example, we measured the bandwidth of the exact circuit recommended by Hanson [80]. We damped the string by wrapping it in felt to reduce




Figure F.12: More uniform radiation pattern of lamp

the effects of the resonance on the measurement. Then, we actuated the string with a swept-sinusoid, measured the response, and computed the linearized frequency response, as shown in Figure F.13 [23]. The only partially damped resonances can be seen as small peaks starting at around 100Hz and moving upward. The bandwidth is clearly limited to the order of about 1kHz. This bandwidth limit was sufficient for Hanson's application of measuring string vibration, but it is not sufficient for our control application.

Hanson's phototransistor type sensors have further limitations in that they are actually designed for digital applications, such as detecting the presence or absence of paper in a printer. Consequently, they do not behave particularly linearly. One problem is that the sensitive area is in the shape of a circle, for which no affine relationship exists between string displacement and output voltage. Furthermore, due to geometrical considerations, for instance the curvature of the photodiode surface as well as likely nonuniform sensitivity, the surface cannot be easily masked to improve linearity. Consequently, these sensors can only be used for making linear measurements when




Figure F.13: Piezoelectric guitar string termination sensor from Graphtech

placed very closely to the end of the string where displacements are small, further limiting their application. Finally, these sensors only work well with very particular diameters of strings, and the sensors have suboptimal noise performance.

F.3.2 Piezoelectric Sensors

Some piezoelectric sensors are mass produced for the application of sensing guitar string vibrations. A piezoelectric sensor is made of an element that forms a charge across itself when it is stressed [8]. When a piezoelectric element is placed directly under a vibrating string termination, it primarily measures the force that the string exerts on it in the vertical dimension due to transverse waves.⁹ The piezoelectric element is typically not connected to any part of the midsection of the string as it could adversely alter the way in which the string vibrates.

Given that such piezoelectric elements typically have small mass and large stiffness, their mechanical resonance frequency is above the hearing threshold [8]. Hence, the force exerted on the piezoelectric element is proportional to the string's vertical transverse displacement very near the end of the string. It follows that the charge on the piezoelectric element is proportional to the vertical string displacement.

Piezoelectric crystal and piezoceramic element manufacturers recommend making measurements using charge amplifiers, which convert the charge on the piezoelectric element to a voltage [8]. A piezoelectric guitar string termination sensor from Graph-tech for a Fender electric guitar is shown in Figure F.14. We have used this type of sensor for the actively controllable electric guitar detailed in Section 4.7.

CNMAT has carried out an experiment on multi-axis string sensing using piezoelectric elements. The string termination is attached to a nearly rigid plate, which flexes according to string vibrations. By forming linear combinations of the signals from the four piezoelectric elements on the underside of the plate, the horizontal and vertical transverse waves as well as longitudinal waves can be measured [66].

⁹For a stringed instrument, we consider the vertical plane to be the plane perpendicular to the top plate. We consider transverse vibrations not orthogonal to the top plate to be horizontal plane vibrations.



Figure F.14: Piezoelectric guitar string termination sensor from Graphtech

F.3.3 Sensor THD Estimation

Here we explain how we estimate the THD of a vibrating string sensor. The technique actually measures the net THD due to the actuator, the nonlinear vibrating string dynamics, and the sensor, but we minimize the actuator and string nonlinearity so that the sensor is the dominant contributing source of nonlinearity. First, we actuate the vibrating string at its fundamental frequency using the linear actuator described in Section F.2.1. Then, we ensure that we vibrate the string at small enough amplitudes that it is likely behaving primarily linearly. In particular, we ensure that we do not observe the resonance shifting effects due to tension modulation as described by Hanson [81]. Thus, we believe that the harmonic distortion observed at the sensor is primarily due to sensor nonlinearity.

We display the sensor input signal on a spectrum analyzer. We follow an example with slightly inferior performance so that the frequency analyzer display is easier to interpret.¹⁰ We estimate the performance of the S2387-1010R silicon photodiode from Hamamatsu illuminated by the halogen lamp. The string was vibrating at a peak-to-peak amplitude of about 3mm at the sensor position.

From Figure F.15, we estimate that the levels of the second, third, fourth, and fifth harmonics are -35dB, -39dB, -42dB, and -49dB lower than the fundamental (first harmonic). Since the higher partials are even smaller in magnitude, we estimate the THD using only these partials:

¹⁰For the preferred configuration with the PD638B photodiode and IR204/H16/L10 LED in Section F.3.1, which was used for experiments in Chapters 3 and 4, we estimate that $k_{displacement} \approx -40\text{dB}$.




Figure F.15: Spectrum analyzer showing the nonlinearity in the sensor signal for a S2387-1010R silicon photodiode illuminated by the halogen lamp

$$k_{displacement} \approx 10 \log_{10}(10^{-35/10} + 10^{-39/10} + 10^{-42/10} + 10^{-49/10}) \approx -33\text{dB}. \quad (\text{F.13})$$

To obtain the string velocity from the string displacement, we differentiate the sensor signal. The differentiator has a magnitude response of +6dB/octave, so to estimate $k_{velocity}$, the THD of the estimated velocity signal, we increase the magnitude of the higher partials by 6dB/octave, as in (F.14).

$$k_{velocity} \approx 10 \log_{10}(10^{(-35+6)/10} + 10^{(-39+6 \log_2 \frac{3}{2})/10} + 10^{(-42+12)/10} + 10^{(-49+6 \log_2 \frac{5}{2})/10}) \quad (\text{F.14})$$

We see that the THD of the velocity signal is even larger than the THD of the displacement signal.

$$k_{velocity} \approx -26\text{dB} \quad (\text{F.15})$$

F.4 Collocation Considerations

As explained in Section 2.4, we need to be able to collocate sensors and actuators, and the sensors and actuators should be matched, if possible. We have chosen to collocate a photodiode optical displacement sensor (see Section F.3.1) with a Lorentz' Law type actuator (see Section F.2.1). We collocate the elements in the presence of the string as shown in Figure F.16. The red semicircle represents the infrared LED, the short red line segment represents the sensitive surface of the photodiode, the long blue line segments represent the magnets, and the thin black line along the middle represents the string (see Figure F.16). Section 2.5 explains why we purposely chose to offset the center of the sensor from the center of the magnets.

When a current I flows along the string, this induces a magnetic field encircling the string. Despite our efforts in shielding the photodiode circuit (see Figure F.8), this magnetic field inevitably cuts through the loop formed by the leads of the photodiode and the operational amplifier. Hence, there exists a string displacement-dependent mutual inductance $L_{as}(x)$ between the actuator and the sensor. This inductance




Figure F.16: Collocation geometry

implies that the current I flowing along the string induces a voltage V_{as} at the input to the operational amplifier [91].

$$V_{as} = L_{as}(x) \frac{\partial I}{\partial t} \quad (\text{F.16})$$

Because of the differentiation in (F.16), the inductive coupling becomes a problem only at higher frequencies. We considered constructing a circuit for directly subtracting out the interference at the input to the operational amplifier. However, we determined that L_{as} varies by about 2.5dB for string displacements 2mm in peak-to-peak amplitude, which would have made canceling the coupling more difficult. In experiments where we needed to reduce the relative influence of the inductive coupling, we simply used the halogen lamp instead of the photodiode as a light source, as this greatly increased the gain between the string displacement and the sensed voltage.

F.5 Active String Termination Concept

The goal of this thesis is to study applications of control theory to musical acoustics, so we have only developed sensors and actuators as necessary. During this process, we became aware of a method for constructing a much more efficient string actuator, but




Figure F.17: String termination actuator concept

we did not pursue the design further. The general concept behind the active string termination is shown in Figure F.17. The string is terminated atop a neodymium magnet, which is held in place by springs. By applying an external magnetic field, a relatively strong force can be exerted on the string.

Appendix G

Toolbox for the Feedback Control of Sound

G.1 Introduction

Feedback control of sound has applications in altering the dynamics of ducts, bars, plates, vibrating strings, and headphones [11]. Analog controllers are fast, but digital controllers are more modern and often allow for simpler and more precise tuning of controller gains. Consider the configuration shown in Figure G.1. Generally the goal is to implement a feedback control transfer function of $K(s)$; however, there is an inherent delay of D seconds due to the digital controller [7].




Figure G.1: Feedback control of a vibrating string (controller model includes a delay of D seconds)

In order to implement the desired controller, D must be small. In particular, to apply classical feedback control to control a vibration at f Hz, we need a controller with system delay $D \ll \frac{1}{f}$ [65]. For example, for $f = 5\text{kHz}$, $D \ll 200\mu\text{s}$.

There is a further bandwidth consideration. The range of human hearing spans roughly from 20Hz to 20kHz [59]. According to the Nyquist-Shannon sampling theorem, the sampling rate, also known as the servo rate, must be at least 40kHz so the whole bandwidth that humans hear can be sampled and reconstructed within the feedback loop [148].

Digital feedback controllers for controlling sound must both be able to process relatively large signal bandwidths in comparison with other control applications, and the delay D (also known as the latency) must still remain small. In many hardware designs, these two requirements are often mutually exclusive [7].

Standard low-latency digital feedback controllers for robotics applications typically do not implement floating point natively, which makes development of algorithms appropriate for the feedback control of sound more difficult [57]. The resulting extra development time is undesirable in research contexts. In addition, many of these robotics controllers cannot operate at high enough sampling rates.

General purpose computers and sound interfaces appear to present a viable alternative, as they can process signals with relatively large audio bandwidths. However, the system delay D is usually prohibitively long. This is due to efficiency considerations in the operating system scheduler as well as in the implementation of anti-aliasing and anti-imaging filters in sigma delta modulators [7]. We measured the minimum delay achievable using a sound interface on the machine described below. Planet CCRMA's low-latency kernel patch made it possible to achieve $D \approx 4\text{ms}$ at a sampling rate of $f_S = 96\text{kHz}$, which would limit classical feedback control to frequencies far below 250Hz [104].

Another alternative we might consider is constructing a low-latency DSP system with a processor supporting native floating point calculations.¹ The most practical configuration for computer music applications involves interfacing a Texas Instruments C6713 DSK with evaluation boards to provide the fast converters.

¹These systems are available commercially, but they are prohibitively expensive.

These boards could for example consist of the 5-6K interface board, the DAC7554 Evaluation Module, and the ADS8361 Evaluation Module. However, we never tested this module because we decided that a non-proprietary and open source alternative would be preferable.

G.2 Open Source Solution

Here we provide details on our solution. It was known that some solution was possible using RTAI, but there was no record anywhere of which hardware and software would be most compatible, or precisely what the performance would be. We hope that this information will be useful to other researchers considering studying the feedback control of sound. We are also releasing all of the source code that we have written.²

The TFCS system consists of the following hardware and software items:

- General purpose computer (AMD Athlon 64 X2 Dual Core 4400+, 1024kb cache)
- Data acquisition card (NI PCI6221 for \$476)
- UNIX-based operating system (Linux Fedora Core 6, Kernel 2.6.19)
- Real-Time Application Interface (RTAI ver 3.5)
- Comedi (Comedi ver 0.7.75 and Comedilib 0.8.1)
- Open Sound Control (OSC from oscpack)
- Pure Data (Pd)
- Synthesis ToolKit (STK ver 4.3.1)
- Control loop code and GUI

²<http://ccrma.stanford.edu/~eberdahl/Projects/TFCS>

To install RTAI, first recompile the kernel with RTAI support. Then install the Comedi data acquisition drivers. Finally, install the TFCS software and reserve one processor or core for running only the control loop [19]. For example, to run the control loop on CPU 1 and all other code on CPU 0, make the following changes:

- Add the kernel switch `isolcpus=1` to `grub.conf`.
- Add the switch `IsolCpusMask=2` when inserting the module `rtai_hal.ko`.
- When creating the real time task in user mode (the control loop), pass `cpus_allowed=2` to `rt_task_init_schmod()`.

G.3 Operation

TFCS has been designed to take full advantage of RTAI. The control code thread is spawned from a user space soft real time process. Then the thread's memory is locked, and the process is made hard real time using RTAI. This trick allows the thread to enable and disable interrupts, which is necessary for hard real time performance [57] (see the dashed lines in Figure G.2).

The graphical user interface (GUI) communicates with the control loop thread using open sound control (OSC) [171]. The controller parameters can be changed in real time from the GUI. For instance, to implement tremolo, a slowly-varying oscillator in the GUI could send OSC messages to the control loop thread, varying the damping parameter over time.

Because OSC is a network protocol, any OSC server can be used to adjust the controller parameters. This feature makes the design attractive for computer music applications because a wide array of software is capable of serving OSC messages.

The GUI for TFCS is implemented using Pure Data (Pd) [131]. The default GUI for TFCS is shown in Figure G.3. It provides a number of elements, such as buttons, toggles, and sliders, for conveniently adjusting the control loop parameters. It is even possible to edit nonlinear characteristic curves in real time using the mouse (see Figure G.4). This characteristic makes it possible to implement a wide array of nonlinear dashpots, nonlinear springs, etc.

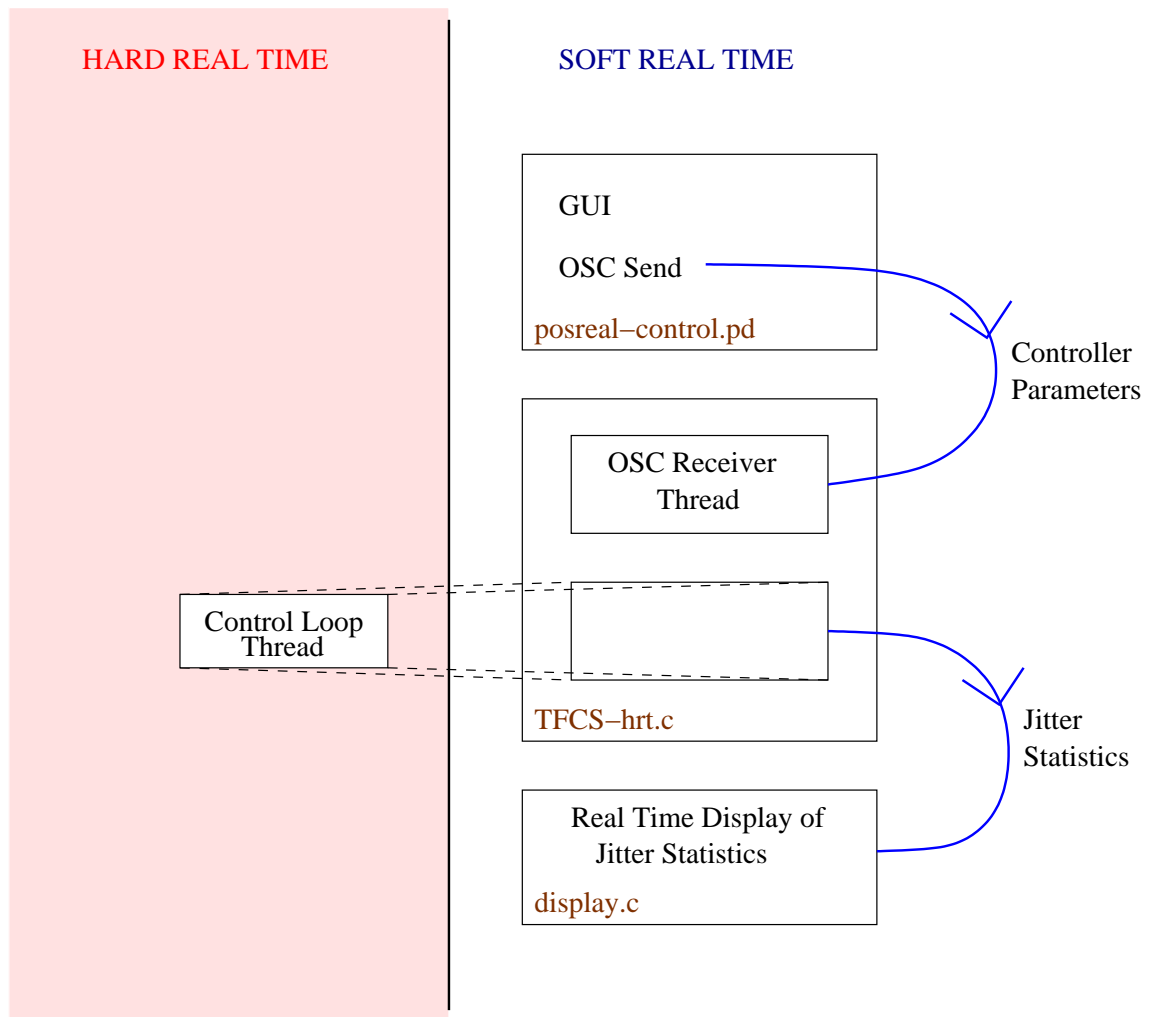


Figure G.2: Modular design of TFCS's software components

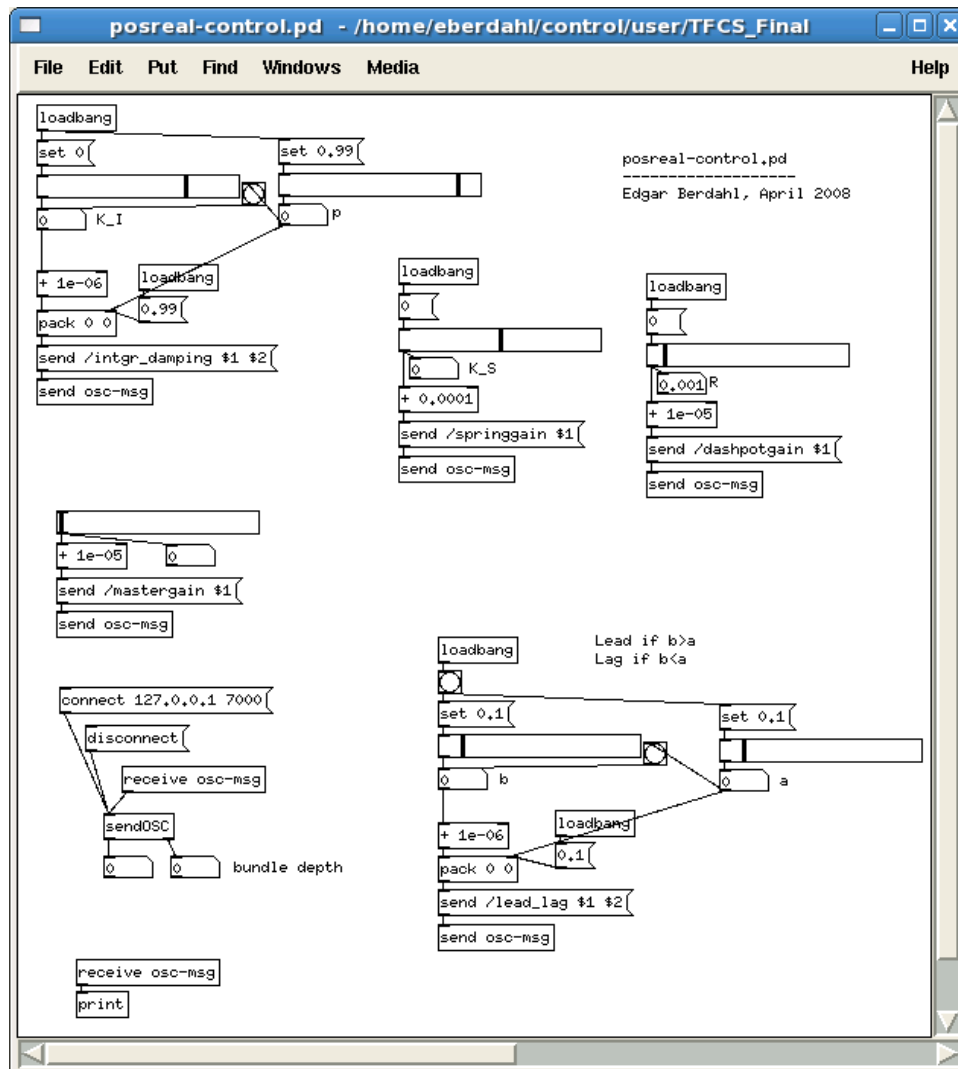


Figure G.3: The GUI for TFCS provides buttons, toggles, sliders, etc.




Figure G.4: The GUI also allows arbitrary nonlinear characteristic curves for controller elements to be drawn with the mouse.




Figure G.5: A continuous time signal (in blue) as approximated by a zero-order hold (dash-dotted in red)

Even if the analog-to-digital converter (ADC), computer, and digital-to-analog converter (DAC) were infinitely fast, the finite sampling rate f_S is nevertheless responsible for a delay. In control engineering, a zero-order hold (ZOH) with no lowpass filter is often used for reconstruction of an analog control signal [65]. Figure G.3 illustrates how sampling is responsible for a delay of half of a sample.

It follows that the total system delay D can be expressed as the sum of the sampling delay and the processing delay $D_{processing}$ due to the finite speed of the ADC, the computer, and the DAC.

$$D = \frac{1}{2f_S} + D_{processing} \quad (\text{G.1})$$

We can estimate D by making a measurement where the controller is configured to simply sample the input signal and then write it to the DAC. Given a sinusoidal input signal, the output signal can be viewed on a scope in a sense analogous to Figure G.3, and the delay can be estimated by visual inspection. For our system, we estimate that $D \approx 24\mu\text{s}$ for $f_S = 40\text{kHz}$. By applying (G.1), we find that $D_{\text{processing}} \approx 11.5\mu\text{s}$.³

G.4 Estimated PDFs Revealing Jitter Performance

While RTAI tries to ensure that the performance is as close to hard real time as is reasonably possible, there is some jitter, or variation in D over time. For instance, the control loop generally communicates with the ADC and DAC via a bus, such as the PCI bus. The graphics card may also use the same PCI bus. If the graphics card and control loop are competing for bus usage, the jitter in D is likely to increase (see Figure G.8).

We present some estimations of the jitter observed when using TFCS under various conditions given the hardware listed in Section G.2. TFCS provides real time feedback about the jitter performance.⁴ In Figure G.6, we present the performance of TFCS under fairly optimal conditions: $f_S = 40\text{kHz}$ and no system loading by other applications. Figure G.6 (top) shows an estimate of how often the control loop runs early or late by various approximate intervals. This estimate is in the form of a discretized estimated probability density function (PDF). Fewer than 4% of samples have jitter larger than 60ns, although we might have been able to develop a tighter bound if we had changed how the PDF estimate was sampled. None of the samples had jitter larger than $1.6\mu\text{s}$ for a test carried out over several minutes. Because so many of the frequencies are so small, Figure G.6 bottom shows the log of the percentage frequencies. Since $\log(0) = -\infty$, no circles are drawn for time intervals during which the control loop was never called (see Figure G.6, bottom).

³For our National Instruments PCI6221 data acquisition card, acquisition takes about $7\mu\text{s}$, while the DAC has a full-scale settling time of roughly $4\mu\text{s}$. The approximate remaining delay may be due to various sources such as acquiring the PCI bus [57].

⁴This feature could be disabled to make it possible to do more computations during the control loop.




Figure G.6: Estimated jitter for TFCS given no additional load ($f_s = 40\text{kHz}$)

Even though the control code is isolated on CPU 0, if CPU 1 is heavily loaded, the jitter performance is affected. Figure G.7 shows the statistics gathered over several minutes. The conditions were the same as for Figure G.6 except that MATLAB was continuously inverting random 400×400 matrices in the background. We see that the jitter was still almost always less than $0.5\mu\text{s}$.

Loading the PCI bus can increase jitter. We believe that our graphics card configuration may have been suboptimal. The statistics shown in Figure G.8 correspond to a test under the same conditions as in Figure G.6 with one exception: the user was constantly dragging an X11 window back and forth. There was also a problem with overruns, which is when the control loop is called so late that it is already time to process the next sample. In this test, about 0.8% of the samples were dropped due to overruns. We found that disabling hardware acceleration for the graphics card worsened the problem.

Finally we tested whether we could increase the sampling rate further. We found $f_s = 50\text{kHz}$ to roughly be the maximum rate (see Figure G.9). Here we were not running the full TFCS control loop, rather the control loop merely read an input




Figure G.7: Estimated jitter for TFCS given significant system load ($f_S = 40kHz$)




Figure G.8: Estimated jitter for TFCS given a system loaded by constantly dragging a window back and forth ($f_S = 40kHz$, 0.8% overruns)




Figure G.9: Estimated jitter for TFCS with no load ($f_S = 50kHz$, 4.6% overruns)

sample, wrote an output sample, updated the jitter statistics, and occasionally passed the jitter statistics to the soft real time display process. We see that with 4.6% overruns, it would probably be better to run at the slower rate of $f_S = 40kHz$. However, we never finished completely optimizing our system. Perhaps by replacing the graphics card we might have obtained better performance.

G.5 Conclusion

The following points explain why we created the Toolbox for the Feedback Control of Sound (TFCS):

1. Implementation is convenient and inexpensive.
2. We present an example configuration.
3. Low-latency high bandwidth control should be available to everyone!
4. Controller parameters can be adjusted over OSC (including over a network).

5. During the course of implementing our own system, we too often came across messages on newsgroups like: "Well, it isn't straight-forward, but we had to write our own code because we couldn't find appropriate starter code anywhere. You will figure it out eventually!"

We wish to thank all of those who helped us set up the TFCS system and investigate alternatives:

- Carr Wilkerson, Fernando Lopez-Lezcano, Chris Chafe, Peter Lindener, and David Yeh at CCRMA
- Heiko Zeuner from Sennheiser Electronics
- Benjamin Faber from Faber Acoustical
- Paolo Mantegazza from the Dipartimento di Ingegneria Aerospaziale, Politecnico di Milano

G.6 Initialization Files

G.6.1 `grub.conf`

The entry for the recompiled kernel in `grub.conf` should be similar to the following:

```
title Fedora Core (2.6.19, By default code runs only on CPU 0)
    root (hd0,0)
    kernel /vmlinuz-2.6.19 ro root=/dev ...
        /VolGroup00/LogVol00 isolcpus=1 noirqbalance
    initrd /initrd-2.6.19.img
```

G.6.2 `/usr/realtime/init_rtai`

Often despite the installation of RTAI, the device nodes need to be recreated after every reboot.

```

export PATH=$PATH:/usr/realtime/bin

if test \! -c /dev/rtai_shm; then
    mknod -m 666 /dev/rtai_shm c 10 254
fi

for n in `seq 0 9`; do
    f=/dev/rtf$n
    if test \! -c $f; then
        mknod -m 666 $f c 150 $n
    fi
done

```

G.6.3 /usr/realtime/start_rtai

The RTAI modules can be loaded using `insmod`.

```

/sbin/insmod /usr/realtime/modules/rtai_hal.ko ...
    rtai_cpufreq_arg=2210762026 ...
    rtai_apicfreq_arg=12561129 IsolCpusMask=2

/sbin/insmod /usr/realtime/modules/rtai_sched.ko ...
    SetupTimeTIMER=3352 Latency=3160

/sbin/insmod /usr/realtime/modules/rtai_fifos.ko

```

G.6.4 /usr/local/comedi/load_cmods

The following commands load the relevant Comedi and RTAI modules for TFCS and our particular data acquisition card (Native Instruments PCI6221). If you don't load the modules in the right order, you may receive a cryptic error message about unrecognized symbols.

```

/sbin/modprobe ni_pcimio

```

```

/usr/local/sbin/comedi_config /dev/comedi0 ni_pcimio
/sbin/modprobe kcomedilib
/sbin/insmod /usr/realtime/modules/rtai_shm.ko
/sbin/insmod /usr/realtime/modules/rtai_msg.ko
/sbin/insmod /usr/realtime/modules/rtai_sem.ko
/sbin/insmod /usr/realtime/modules/rtai_mbx.ko
/sbin/insmod /usr/realtime/modules/rtai_comedi.ko

```

G.6.5 kcomedilib.c missing functionality

In `my_kcomedilib.c`, we implemented some helper functions for getting information on the subdevice to be used as well as for converting between quantities the subdevice understands and physical quantities. These functions behave analogously to the functions in the `comedilib` library, which is currently not directly compatible with RTAI.

```

my_comedi_subdev *get_subdev_info(void *dev,
    int subdev_code, unsigned int range,
    int channel, char *nm)

double inline my_comedi_to_phys(lsampl_t data,
    my_comedi_subdev *subdev)

lsampl_t inline my_comedi_from_phys(double data,
    my_comedi_subdev *subdev)

```

Bibliography

- [1] Jonathan Abel and Dave Berners. On peak-detecting and rms feedback and feedforward compressors. In *Proceedings of the 115th Convention of the Audio Engineering Society*, Preprint #5914, 2003.
- [2] Richard Adams and Blake Hannaford. Stable haptic interaction with virtual environments. *IEEE Transactions on Robotics and Automation*, 15(3):465–474, June 1999.
- [3] William Aitken, Anthony Sedivy, and Michael Dixon. Electronic musical instrument. International Patent, WO/1984/004619, 1984.
- [4] Thomas Alberts and Joseph Colvin. Observations on the nature of transfer functions for control of piezoelectric laminates. *Journal of Intelligent Material Systems and Structures*, 2:528–541, 1991.
- [5] Mehdi Ammi and Antoine Ferreira. Robotic assisted micromanipulation system using virtual fixtures and metaphors. In *Proc. IEEE International Conference on Robotics and Automation*, pages 454–460, Rome, Italy, April 2007.
- [6] Eric Anderson and Nesbitt Hagood. Simultaneous piezoelectric sensing/actuation: Analysis and application to controlled structures. *Journal of Sound and Vibration*, 174(5):617–639, 1994.
- [7] Markko Antila. Contemporary electronics solutions for active noise control. In *Proceedings of the International Symposium on Active Noise and Vibration Control*, Williamsburg, VA, September 2004.

- [8] Antonio Arnau. *Piezoelectric transducers and applications*. Springer, New York, NY, 2004.
- [9] Ralf Austermann. Circuit arrangement for commutating a reluctance motor. U.S. Pat. No. 5,180,960, January 1993.
- [10] Mark Balas. Direct velocity feedback control of large space structures. *Journal of Guidance and Control*, 2:252–253, 1979.
- [11] Volker Bartels. Headset with active noise-reduction system for mobile applications. *Journal of the Audio Engineering Society*, 40(4):277–281, 1992.
- [12] R.C. Becerra, T.M. Jahns, and Mehrdad Ehsani. Four-quadrant sensorless brushless ECM drive. In *Proceedings of the Sixth Annual Applied Power Electronics Conference and Exposition*, pages 202–209, Dallas, TX, March 1991.
- [13] Peter Bennett. Hap-kit: A haptic interface for a virtual drum-kit. In *The Symposium for Cybernetics Annual Research Projects*, University of Reading, June 2004.
- [14] Edgar Berdahl. HSP: A simple and effective open-source platform for implementing haptic musical instruments, Last checked: July 27, 2009, <http://ccrma.stanford.edu/~eberdahl/Projects/HSP>.
- [15] Edgar Berdahl. Lab 6 for Music 250—Physical Interaction Design for Music, Last checked: July 27, 2009, <http://ccrma.stanford.edu/courses/250a/labs/Hapticslab>.
- [16] Edgar Berdahl, Steven Backer, and Julius Smith. If I Had A Hammer: Design and theory of an electromagnetically-prepared piano. In *Proceedings of the International Computer Music Conference*, pages 81–84, Barcelona, Spain, September 2005. Computer Music Association.
- [17] Edgar Berdahl and Julius O. Smith III. Some physical audio effects. In *Proceedings of the 9th International Conference on Digital Audio Effects*, pages 165–168, Sept. 18–20 2006.

- [18] Edgar Berdahl and Julius O. Smith III. Estimating the state of a one-dimensional waveguide. 153rd Meeting of the Acoustical Society of America, June 4-7 2007.
- [19] Edgar Berdahl, Nelson Lee, Günter Niemeyer, and Julius O. Smith III. TFCS: Toolbox for the feedback control of sound. Acoustics '08: 155th Meeting of the Acoustical Society of America, 5th FORUM ACUSTICUM, and 9th Congr s Francais d'Acoustique, June 29-July 4 2008.
- [20] Edgar Berdahl, Günter Niemeyer, and Julius O. Smith III. Applications of passivity theory to the active control of acoustic musical instruments. In *Acoustics '08: 155th Meeting of the Acoustical Society of America, 5th FORUM ACUSTICUM, and 9th Congr s Francais d'Acoustique*, June 2008.
- [21] Edgar Berdahl, Günter Niemeyer, and Julius O. Smith III. Feedback control of acoustic musical instruments. Technical report STAN-M-120, Stanford University, June 2008.
- [22] Edgar Berdahl, Günter Niemeyer, and Julius O. Smith. Using haptic devices to interface directly with digital waveguide-based musical instruments. In *Proceedings of the Ninth International Conference on New Interfaces for Musical Expression*, pages 183–186, Pittsburgh, PA, June 4-6 2009.
- [23] Edgar Berdahl and Julius Smith. Transfer function measurement toolbox, Last checked: July 27, 2009, http://ccrma.stanford.edu/realsimple/imp_meas/.
- [24] Edgar Berdahl and Julius Smith. Active damping of a vibrating string. In *Proceedings of the International Symposium on Active Noise and Vibration Control*, Sept. 18-20 2006.
- [25] Edgar Berdahl and Julius O. Smith. PID control of a vibrating string. In *153rd Meeting of the Acoustical Society of America*, June 4-7 2007.

- [26] Edgar Berdahl, Hans-Christoph Steiner, and Collin Oldham. Practical hardware and algorithms for creating haptic musical instruments. In *Proceedings of the International Conference on New Interfaces for Musical Expression (NIME-2008)*, Genova, Italy, 2008.
- [27] Edgar Berdahl, Bill Verplank, Julius O. Smith, and Günter Niemeyer. A physically intuitive haptic drumstick. In *Proceedings of the International Computer Music Conference*, pages 363–6, Denmark, Copenhagen, August 2007.
- [28] Charles Besnainou. Transforming the voice of musical instruments by active control of the sound radiation. In *Proceedings of the International Symposium on Active Noise and Vibration Control*, Fort Lauderdale, FL, 1999.
- [29] Jeff Bilmes. Techniques to foster drum machine expressivity. In *Proceedings of the International Computer Music Conference*, Japan, 1993.
- [30] David Birnbaum. Musical vibrotactile feedback. Master’s thesis, McGill University, Montreal, Canada, October 2007.
- [31] David Birnbaum, Rebecca Fiebrink, J. Malloch, and Marcelo Wanderley. Towards a dimension space for musical devices. In *Proceedings International Conference on New Interfaces for Musical Expression (NIME05)*, pages 192–195, Vancouver, BC, Canada, 2005.
- [32] Max Born and Emil Wolf. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, New York, NY, 7th ed. edition, 1999.
- [33] Henri Boutin. Contrôle actif sur instruments acoustiques. ATIAM Master’s Thesis, Laboratoire d’Acoustique Musicale, Université Pierre et Marie Curie, Sept. 2007.
- [34] Henri Boutin and Charles Besnainou. Physical parameters of an oscillator changed by active control: Application to a xylophone bar. In *Proceedings*

- of the 11th International Conference on Digital Audio Effects*, Espoo, Finland, Sept. 1-4 2008.
- [35] Henri Boutin and Charles Besnainou. Physical parameters of the violin bridge changed by active control. In *Acoustics '08: 155th Meeting of the Acoustical Society of America, 5th FORUM ACUSTICUM, and 9th Congr  s Francais d'Acoustique*, June, 2008.
- [36] Morten Bro-Nielsen. Finite element modeling in surgery simulation. *Proceedings of the IEEE*, 86(3):490–503, March 1998.
- [37] Cynthia Bruyns. Modal synthesis for arbitrarily shaped objects. *Computer Music Journal*, 30(3):22–37, 2006.
- [38] Bryan and Harter. Studies in the physiology and psychology of the telegraphic language. *Psychological Review*, 4:27–53, 1897.
- [39] Juan-Pablo Caceres, Robert Hamilton, Deepak Iyer, Chris Chafe, and Ge Wang. To the edge with China: Explorations in network performance. In *ARTECH 2008: Proceedings of the 4th International Conference on Digital Arts*, pages 61–66, Porto, Portugal, 2008.
- [40] Claude Cadoz. Retour au r  el: Le sens du feedback, r  troaction-interaction-enaction. In *Proc. Le feedback dans la cr  ation musicale: Rencontres musicales pluridisciplinaires*, Lyon, France, 2006.
- [41] Claude Cadoz, Annie Luciani, and Jean-Loup Florens. Synth  se musicale par simulation des m  canismes instrumentaux. *Revue d'acoustique*, 59:279–292, 1981.
- [42] Claude Cadoz, Annie Luciani, and Jean-Loup Florens. CORDIS-ANIMA: A modeling and simulation system for sound and image synthesis—The general formalism. *Computer Music Journal*, 17(1):19–29, Spring 1993.
- [43] Graham Carey. Parallelism in finite element modeling. In *Communications in Applied Numerical Methods*, volume 2, pages 281–287, 1986.

- [44] Jim Carter and David Fourney. Research based tactile and haptic interaction guidelines. In *Proceedings of GOTH1'05 Guidelines On Tactile and Haptic Interactions Conference*, volume pp. 84-92, Saskatoon, Saskatchewan, Canada, October 2005.
- [45] Nicolas Castagne and Claude Cadoz. Creating music by means of ‘physical thinking’: The musician oriented Genesis environment. In *Proc. 5th Internat’l Conference on Digital Audio Effects*, pages 169–174, Hamburg, Germany, Sept. 2002.
- [46] Nicolas Castagne, Claude Cadoz, Jean-Loup Florens, and Annie Luciani. Haptics in computer music: A paradigm shift. In *Proceedings of EuroHaptics*, pages 422–425, Munich, Germany, June 5-7 2004.
- [47] Chris Chafe. Tactile audio feedback. In *Proc. Internat’l Computer Music Conf.*, pages 76–79, 1993.
- [48] Remy Chollet, Guillaume Aeberli, and Charles Besnainou. Modifier la resonance de helmholtz de la guitare par contrôle actif. In *Proceedings of the 5th French Congress on Acoustics*, pages 248–250, Sept. 3-6 2000.
- [49] Daniel Cole and Robert Clark. Adaptive compensation of piezoelectric sensori-actuators. *Journal of Intelligent Material Systems and Structures*, 5(5):665–672, 1994.
- [50] Daniela Constantinescu, Septimiu Salcudean, and Elizabeth Croft. Haptic rendering of rigid body contacts using impulsive and penalty forces. *IEEE Transactions on Robotics*, 2005.
- [51] Francois Conti, Federico Barbagli, Dan Morris, Chris Sewell, and Sebastien Grange. CHAI 3D, Last checked: July 27, 2009, <http://www.chai3d.org/>.
- [52] Perry Cook. *Identification of Control Parameters in an Articulatory Vocal Tract Model, With Applications to the Synthesis of Singing*. PhD thesis, Stanford University, 1991.

- [53] Perry Cook. Principles for designing computer music controllers. In *Proc. ACM CHI Workshop in New Interfaces for Musical Expression*, pages 1–4, Seattle, Washington, 2001.
- [54] Perry Cook and Gary Scavone. The synthesis toolkit (STK), version 2.1. In *Proc. Internat'l Computer Music Conf.*, Beijing, China, 1999. See <http://ccrma.stanford.edu/software/stk>.
- [55] Egbert De Boer. Theory of motional feedback. *IRE Transactions on Audio*, 9(1):15–21, 1961.
- [56] Nicola Diolaiti, Guenter Niemeyer, Federico Barbagli, Kenneth Salisbury, and Claudio Melchiorri. The effect of quantization and coulomb friction on the stability of haptic rendering. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 237–246, Pisa, Italy, March 18-20 2005.
- [57] Lorenzo Dozio and Paolo Mantegazza. General-purpose processors for active vibro-acoustic control: Discussion and experiences. *Control Engineering Practice*, 15(2):163–176, February 2007.
- [58] Wang Enping and Wang Chaozhu. Some properties of the positive real matrix. *Acta Mathematicae Applicatae Sinica*, 1(1), 1984.
- [59] Hugo Fastl and Eberhard Zwicker. *Psychoacoustics: Facts and Models*. Springer, New York, NY, 3rd edition, December 2006.
- [60] Jeff Feasel. A virtual stringed instrument with haptic feedback, Last checked: July 27, 2009, <http://www.cs.unc.edu/~feasel/classes/259/final/progress.html>.
- [61] David Feygin, Madeleine Keehner, and Frank Tendick. Haptic guidance: experimental evaluation of a haptic training method for a perceptual motor skill. In *Proceedings of the 10th Symposium On Haptic Interfaces For Virtual Environment and Teleoperator Systems*, pages 40–47, Orlando, FL, March 24-25 2002.

- [62] Neville Fletcher and Thomas Rossing. *The Physics of Musical Instruments*. Springer, New York, NY, 2nd edition, 1998.
- [63] Jean-Loup Florens. Real-time bowed string synthesis with force feedback gesture interaction. In *Proc. of the Forum Acusticum*, Sevilla, Spain, 2002.
- [64] Jesse Fox. Haptic plectrum, Last checked: July 27, 2009, <http://ccrma.stanford.edu/~jrobfox/220c/index.htm>.
- [65] Gene Franklin, J. Powell, and Abbas Emami-Naeini. *Feedback Control of Dynamic Systems*. Prentice Hall, Fifth edition, 2005.
- [66] Adrian Freed and Osman Isvan. Musical applications of new, multi-axis guitar string sensors. In *Proc. of the Int. Computer Music Conf.*, Aug. 27-Sept. 1, 2000.
- [67] Michael Friswell and J. E. Mottershead. *Finite Element Model Updating in Structural Dynamics*. Kluwer Academic Publishers, Boston, MA, 1995.
- [68] Brent Gillespie, Sile O'Modhrain, Philip Tang, David Zaretsky, and Cuong Pham. The virtual teacher. In *Proceedings of the ASME International Mechanical Engineering Congress*, Anaheim, California, November 15-20 1998.
- [69] Jerald Graeme. *Photodiode Amplifiers: OP AMP Solutions*. McGraw-Hill Professional, Columbus, OH, 1995.
- [70] Graph Tech. Company website, Last checked: July 27, 2009, <http://www.graphtech.com/>.
- [71] Steve Griffin, Steven Lane, and Robert Clark. The application of smart structures toward feedback suppression in amplified acoustic guitars. *Journal of the Acoustical Society of America*, 113(6):3188–3196, 2003.
- [72] Graham Grindlay. Haptic guidance benefits musical motor learning. In *Proc. of Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 397–404, Reno, NV, March 2008.

- [73] Jean Guérard and Xavier Boutillon. Real time acoustic wave separation in a tube. In *Proceedings of the International Symposium on Musical Acoustics*, pages 451–456, Edinburgh, Scotland, August 19–22 1997.
- [74] Gary Guthart and Kenneth Salisbury. The intuitive telesurgery system: Overview and application. In *Proc. IEEE International Conference on Robotics and Automation*, volume 1, pages 618–621, San Francisco, CA, April 24–28 2000.
- [75] Nesbitt Hagood and Eric Anderson. Simultaneous sensing and actuation using piezoelectric materials. *Proceedings of the SPIE*, 1543:409–421, 1992.
- [76] Jürgen Hahn, Thomas Edison, and Thomas Edgar. A note on stability analysis using bode plots. *Chemical Engineering Education*, 35(3):208–211, 2001.
- [77] Aram Hajian, Daniel Sanchez, and Robert Howe. Drum roll: Increasing bandwidth through passive impedance modulation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2294–9, Albuquerque, New Mexico, April 20–25 1997.
- [78] Sathya Hanagud and Steven Griffin. Active structural control for a smart guitar. In *Proceedings of the Fourth European Conference On Smart Structures and Materials*, Harrogate, U.K., July 6–8 1998.
- [79] Blake Hannaford. A design framework for teleoperators with kinesthetic feedback. *IEEE Transactions on Robotics and Automation*, 5(4):426–434, August 1989.
- [80] Roger Hanson. Optoelectronic detection of string vibration. In *The Physics Teacher*, volume 25, pages 165–166, 1987.
- [81] Roger Hanson, James Anderson, and Kent Macomber. Measurements of non-linear effects in a driven vibrating wire. *Journal of the Acoustical Society of America*, 96(3):1549–1556, September 1994.
- [82] Gregory Heet. String instrument vibration initiator and sustainer. U.S. Pat. 4,075,921, 1978.

- [83] Cyrille Henry. pmpd: Physical modelling for pure data. In *Proc. Internat'l Computer Music Conf.*, Miami, FL, 2004. See <http://drpichon.free.fr/pmpd/>.
- [84] Alan Hoover. Controls for musical instrument sustainers. U.S. Pat. 6,034,316, 2000.
- [85] Bart Hopkin and John Scoville. *Musical Instrument Design: Practical Information for Instrument Design*. See Sharp Press, Tucson, AZ, 1996.
- [86] Felix Huang, Brent Gillespie, and Arthur Kuo. Human adaptation to interaction forces in visuo-motor coordination. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 14(3):390–397, Sept. 2006.
- [87] Amy Hurst, Jennifer Mankoff, Anind Dey, and Scott Hudson. Dirty desktops: Using a patina of magnetic mouse dust to make common interactor targets easier to select. In *Proc. 20th ACM Symposium on User Interface Software and Technology*, pages 183–186, Newport, Rhode Island, October 2007.
- [88] Iqbal Husain and Mehrdad Ehsani. Rotor position sensing in switched reluctance motor drives by measuring mutually induced voltages. *IEEE Transactions on Industry Applications*, 30(3):665–672, 1994.
- [89] Faustina Hwang, Simeon Keates, Patrick Langdon, and P. John Clarkson. Multiple haptic targets for motion-impaired computer users. In *Proc. Conf. Human Factors in Computing Systems*, pages 41–48, Ft. Lauderdale, Florida, 2003.
- [90] Paul Ierymenko. Unitary transducer control system. U.S. Pat. No. 6,216,059, April 2001.
- [91] Umran Inan and Aziz Inan. *Engineering Electromagnetics*. Addison Wesley Longman, Menlo Park, CA, 1999.
- [92] J. Douglas Rollow IV. *Active Control of Spectral Detail Radiated by an Air-Loaded Impacted Membrane*. PhD thesis, Pennsylvania State University, 2003.

- [93] Nikil Jayant and Peter Noll. *Digital Coding of Waveforms: Principles and Applications to Speech and Video*. Prentice Hall, March 1984.
- [94] Marty Johnson and Stephen Elliott. Active control of sound radiation using volume velocity cancellation. *Journal of the Acoustical Society of America*, 98(4):2174–2186, October 1995.
- [95] Matti Karjalainen, Teemu Mäki-Patola, Aki Kanerva, and Antti Huovilainen. Virtual air guitar. *Journal of the Audio Engineering Society*, 54(10):964–980, October 2006.
- [96] Kevin Karplus and Alexander Strong. Digital synthesis of plucked string and drum timbres. *Computer Music Journal*, 7(2):43–55, 1983.
- [97] Loic Kessousand, Julien Castet, and Daniel Arfib. ‘GXtar’, an interface using guitar techniques. In *Proc. of the International Conf. on New Interfaces for Musical Expression*, pages 192–195, 2006.
- [98] Hassan Khalil. *Nonlinear Systems, 3rd Ed.* Prentice Hall, Upper Saddle River, NJ, 2001.
- [99] Wolfgang Klippel. Loudspeaker nonlinearities - Causes, parameters, symptoms. In *Proceedings of the 119th Convention of the Audio Engineering Society* *Audio Engineering Society*, Preprint #6584, New York, NY, October, 7-10 2005.
- [100] Alexandros Kontogeorgakopoulos and Claude Cadoz. Interfacing digital waveguide with CORDIS ANIMA networks. In *Acoustics '08: 155th Meeting of the Acoustical Society of America, 5th FORUM ACUSTICUM, and 9th Congr s Francais d’Acoustique*, Paris, France, June 29-July 4 2008.
- [101] Katherine Kuchenbecker, Jonathan Fiene, and G nter Niemeyer. Improving contact realism through event-based haptic feedback. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):219–230, March/April 2006.
- [102] Nelson Lee and Julius O. Smith. Vibrating-string coupling estimation from recorded tones. In *Acoustics '08: 155th Meeting of the Acoustical Society of*

- America, 5th FORUM ACUSTICUM, and 9th Congr  s Francais d'Acoustique*, Paris, France, June 2008.
- [103] Herv   Lissek and Florian Sandoz. From the electrical shunting of a loudspeaker to active impedance control. In *Acoustics '08: 155th Meeting of the Acoustical Society of America, 5th FORUM ACUSTICUM, and 9th Congr  s Francais d'Acoustique*, Paris, France, June 29 - July 4 2008.
- [104] Fernando Lopez-Lezcano. Surviving on Planet CCRMA: Two years later and still alive. In *Proceedings of the 3rd International Linux Audio Conference*, pages 109–114, Karlsruhe, Germany, April 2005.
- [105] Kyle Machulis. Novint Falcon open source library, Last checked: July 27, 2009, <http://sourceforge.net/projects/libnifalcon>.
- [106] Max Mathews. The digital computer as a musical instrument. *Science*, 142(3592):553–557, November 1963.
- [107] Robert McGill, John Tukey, and Wayne Larsen. Variations of box plots. *The American Statistician*, 32(1):12–16, February 1978.
- [108] M.E. McIntyre, Robert Schumacher, and Jim Woodhouse. On the oscillations of musical instruments. *Journal of the Acoustical Society of America*, 74(5):1325–1345, November 1983.
- [109] Timothy Miller and Robert Zeleznik. The design of 3D haptic widgets. In *Proc. Symposium on Interactive 3D Graphics*, pages 97–102, Atlanta, Georgia, April 1999.
- [110] Eduardo Miranda, Ross Kirk, and Marcelo Wanderley. *New Digital Musical Instruments*. A-R Editions, Middleton, WI, 2006.
- [111] Bill Moggridge. *Designing Interactions*. MIT Press, Cambridge, MA, 2007.
- [112] Dan Morris, Hong Tan, Federico Barbagli, Timothy Chang, and Kenneth Salisbury. Haptic feedback enhances force skill learning. In *Proc. Second Joint*

- EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 21–26, Tsukuba, Japan, March 2007.
- [113] William Moss and Bryan Cunitz. Haptic theremin: Developing a haptic musical controller using the Sensable Phantom Omni. In *Proceedings of the International Computer Music Conference*, Barcelona, Spain, Sept. 2005.
- [114] Georg Mueller and Werner Lauterborn. The bowed string as a nonlinear dynamical system. *Acustica*, 82:657–664, 1996.
- [115] Philip Nelson and Steve Elliott. *Active Control of Sound*. Academic Press, Elsevier Publishing, Burlington, MA, 1993.
- [116] Bruno Nettl. "Music." Grove Music Online Encyclopedia. Oxford Music Online, Last checked: 13 Nov. 2008, <http://www.oxfordmusiconline.com/subscriber/article/grove/music/40476>.
- [117] Ian Oakley. *Haptic Augmentation of the Cursor: Transforming Virtual Actions into Physical Actions*. PhD thesis, University of Glasgow, Glasgow, Scotland, May 2003.
- [118] Ian Oakley, Alison Adams, Stephen Brewster, and Philip Gray. Guidelines for the design of haptic widgets. In *Proc. 16th British HCI Group Annual Conference*, pages 195–211, London, England, Sept. 2002.
- [119] Ian Oakley, Stephen Brewster, and Philip Gray. Solving multi-target haptic problems in menu interaction. In *Proc. Conference on Human Factors in Computing Systems*, pages 357 – 358, Seattle, Washington, 2001.
- [120] Sile O'Modhrain. *Playing By Feel: Incorporating Haptic Feedback into Computer-Based Musical Instruments*. PhD thesis, Stanford University, Stanford, CA, USA, 2000.
- [121] Open Dynamics Engine. Last checked: July 27, 2009, <http://www.ode.org/>.

- [122] Alan Oppenheim and Ronald Schaffer. *Discrete-Time Signal Processing*. Prentice Hall, Upper Saddle River, NJ, 2nd edition, 1999.
- [123] Brad Osgood. *Lecture Notes for EE261: The Fourier Transform and its Applications*. Stanford University, 2007.
- [124] Flavien Picon, Mehdi Ammi, and Patrick Bourdot. Case study of haptic methods for selection on CAD models. In *Proc. IEEE Virtual Reality Conference*, pages 209–212, Reno, NV, March 2008.
- [125] Flavien Picon, Mehdi Ammi, and Patrick Bourdot. Force model for CAD selection. In *Proc. ACM symposium Virtual reality software and technology*, pages 283–284, Bordeaux, France, 2008.
- [126] Allen Pierce. *Acoustics: An Introduction to Its Physical Principles and Applications*. Acoustical Society of America, June 1989.
- [127] Nick Porcaro, David Jaffe, Pat Scandalis, Julius Smith, Tim Stilson, and Scott Van Duyne. Synthbuilder: A graphical rapid-prototyping tool for the development of music synthesis and effects patches on multiple platforms. *Computer Music Journal*, 22(2):35–43, Summer 1998.
- [128] L. Postman and J. Egan. *Experimental psychology: An introduction*. Harper, New York, 1949.
- [129] Francois Poyer and Claude Cadoz. Sound synthesis and musical composition by physical modeling of self-sustained oscillating structures. In *Proc. of 4th Sound and Music Computing Conference*, pages 14–21, Lefkada, Greece, July 2007.
- [130] Aure Prochazka. VG-8.com, Last checked: July 27, 2009, <http://www.vg-8.com/>.
- [131] Miller Puckette. Pure data: another integrated computer music environment. In *Proceedings of the Second Intercollege Computer Music Concerts*, pages 37–41, Tachikawa, Japan, 1997.

- [132] Miller Puckette, Theodore Apel, and David Zicarelli. Real-time audio analysis tools for pd and msp. In *Proc. of the Int. Computer Music Conf.*, pages 109–112, 1998.
- [133] W. Ray and Ibrahim Al-Bahadly. Sensorless methods for determining the rotor position of switched reluctance motors. In *Proceedings of the Fifth European Conference on Power Electronics and Applications*, volume 6, pages 7–13, 1993.
- [134] Curtis Roads. *The Computer Music Tutorial*. MIT Press, Cambridge, MA, February 1996.
- [135] Xavier Rodet and Chriophe Vergez. Nonlinear dynamics in physical models: From basic models to true musical-instrument models. *Computer Music Journal*, 23(3):35–49, 1999.
- [136] Louis Rosenberg and Scott Brave. Using force feedback to enhance human performance in graphical user interfaces. In *Proc. Conference on Human Factors in Computing Systems*, pages 291–292, Vancouver, British Columbia, Canada, April 1996.
- [137] Kenneth Salisbury and Christopher Tarr. Haptic rendering of surfaces defined by implicit functions. In *Proc. of the ASME 6th Annual Symposium on Haptic interfaces for Virtual Environment and Teleoperator Systems*, pages 61–68, Dallas, TX, November 1997.
- [138] Shankar Sastry. *Nonlinear Systems: Analysis, Stability, and Control*. Springer, New York, NY, 2004.
- [139] Richard Schmidt. *Motor Control and Learning: A Behavioral Emphasis*. Human Kinetics Publishers, Champaign, Illinois, 1982.
- [140] Stefania Serafin. *The Sound of Friction: Real-Time Models, Playability and Musical Applications*. PhD thesis, Stanford University, Stanford, CA, USA, June 2004.

- [141] Stefania Serafin, Matthew Burtner, Charles Nichols, and Sile O'Modhrain. Expressive controllers for bowed string physical models. In *Proc. 1st Internat'l Conference on Digital Audio Effects*, Limerick, Ireland, December 2001.
- [142] Reza Shadmehr and Steven Wise. *The Computational Neurobiology of Reaching and Pointing*. MIT Press, Cambridge, MA, 2005.
- [143] Shaffer. Cognition and motor programming. *Tutorials in motor neuroscience*, pages 371–383, 1991.
- [144] C. Sherman and J. Butler. Analysis of harmonic distortion in electroacoustic transducers. *Journal of the Acoustical Society of America*, 98(3):1596–1611, September 1995.
- [145] Stephen Sinclair, Gary Scavone, and Marcelo Wanderley. Audio-haptic interaction with the digital waveguide bowed string. In *Proc. Internat'l Computer Music Conf.*, Montreal, Canada, August 2009.
- [146] Stephen Sinclair and Marcelo Wanderley. A run-time programmable simulator to enable multi-modal interaction with rigid-body systems. *Interacting with Computers*, 21(1-2):54–63, January 2009.
- [147] Julius O. Smith. *Introduction to Digital Filters with Audio Applications*. W3K Publishing, <http://ccrma.stanford.edu/~jos/pasp/>, 2007.
- [148] Julius O. Smith. *Mathematics of the Discrete Fourier Transform (DFT) with Audio Applications*. W3K Publishing, <http://ccrma.stanford.edu/~jos/mdft/>, 2008.
- [149] Julius O. Smith. *Physical Audio Signal Processing: For Virtual Musical Instruments and Audio Effects*. W3K Publishing, <http://ccrma.stanford.edu/~jos/pasp/>, December 2008.
- [150] Starr Labs. Company website, Last checked: July 27, 2009, <http://www.starrlabs.com/>.

- [151] Tim Stilson. Forward-going wave extraction in acoustic tubes. In *Proceedings of the International Computer Music Conference*, Banff, Canada, 1995.
- [152] Karl Åström. Event based control. *Analysis and Design of Nonlinear Control Systems*, pages 127–147, 2008.
- [153] Karl Åström and Bjorn Wittenmark. *Adaptive Control*. Addison-Wesley, Boston, MA, 2nd edition, 1995.
- [154] J. Q. Sun. Some observations on physical duality and colocation of structural control sensors and actuators. *Journal of Sound and Vibration*, 194(5):765–770, 1996.
- [155] Nicolas Szilas and Claude Cadoz. Analysis techniques for physical modeling networks. *Computer Music Journal*, 22(3):33–48, Fall 1998.
- [156] Leon Theremin. Method of and apparatus for the generation of sounds. U.S. Pat. No. 1,661,058, December 1925.
- [157] Ulrich Tietze and Christoph Schenk. *Halbleiter-Schaltungstechnik*. Springer, Berlin, Germany, 11th edition, 1999.
- [158] C. Touzé, O. Thomas, and Antoine Chaigne. Hardening/softening behavior in non-linear oscillations of structural systems using non-linear normal modes. *Journal of Sound and Vibration*, 273:77–101, 2004.
- [159] Dan Trueman. Why a laptop orchestra? *Organised Sound*, 12(2):171–179, 2007.
- [160] Jaine Tsai. Technical data sheet, 3mm infrared led t-1 ir204/h16/l10. Technical report, Everlight Corp., 2005.
- [161] Mac Van Valkenburg. *Introduction to Modern Network Synthesis*. John Wiley and Sons Inc., Hoboken, NJ, 1960.
- [162] Balthasar van der Pol. The nonlinear theory of electric oscillations. In *Proceedings of the Institute of Radio Engineers*, volume 22, pages 1051–1086, 1934.

- [163] Maarten van Walstijn and Pedro Rebelo. The prosthetic conga: Towards an actively controlled hybrid musical instrument. In *Proceedings of the International Computer Music Conference*, pages pp. 786–789, Sept. 5-9 2005.
- [164] Vincent Verfaille, Udo Zölzer, and Daniel Arfib. Adaptive digital audio effects: A new class of sound transformations. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(5), September 2006.
- [165] Bill Verplank. Interaction design sketchbook. Class lecture, Stanford University, Stanford, CA, September 2003. <http://ccrma.stanford.edu/courses/250a/lectures/IDSketchbok.pdf>.
- [166] Bill Verplank. Haptic music exercises. In *Proc. of the Int. Conf. on New Interfaces for Musical Expression*, pages 256–257, Vancouver, BC, Canada, 2005.
- [167] Hermann von Helmholtz. *Die Lehre von den Tonempfindungen als physiologische Grundlage für die Theorie der Musik*. Friedrich Vieweg und Sohn, 1896.
- [168] Gabriel Weinreich and René Caussé. Digital and analog bows: Hybrid mechanical-electrical systems. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 11, pages 1297–1299, 1986.
- [169] Daniel Wolpert and Mitsuo Kawato. Multiple paired forward and inverse models for motor control. *Neural Networks*, 11(7-8):1317–29, Oct.-Nov. 1998.
- [170] Matt Wright and Edgar Berdahl. Toward machine learning of expressive microtiming in brazilian drumming. In *Proceedings of the International Computer Music Conference*, New Orleans, Louisiana, November 6-11 2006.
- [171] Matt Wright and Adrian Freed. Open Sound Control: A new protocol for communicating with sound synthesizers. In *Proceedings of the International Computer Music Conference*, pages 101–104, Thessaloniki, Hellas, September 1997.

- [172] Chai Wah Wu. Qualitative analysis of dynamic circuits. *Wiley Encyclopedia of Electrical and Electronics Engineering*, 1999.
- [173] Toshio Yamada, Daisuke Tsubouchi, Tetsuro Ogi, and Michitaka Hirose. Desk-sized immersive workplace using force feedback grid interface. In *Proc. IEEE Virtual Reality Conference*, pages 135–142, Washington, DC, USA, 2002.
- [174] Yamaha Corporation. Company website, Last checked: July 27, 2009, <http://www.yamaha.com/>.
- [175] Yasuyoshi Yokokohji, Ralph Hollis, Takeo Kanade, Kazuyuki Henmi, and Tsuneo Yoshikawa. Toward machine mediated training of motor skill. In *Proceedings of 1993 IEEE International Symposium on Robot and Human Communication*, pages 32–37, 1996.
- [176] M. Yunik, M. Borys, and G. Swift. A digital flute. *Computer Music Journal*, 9(2):49–52, Summer 1985.