

CENTER FOR COMPUTER RESEARCH IN MUSIC AND ACOUSTICS
DEPARTMENT OF MUSIC, STANFORD UNIVERSITY
REPORT NO. 120

Feedback Control of Acoustic Musical Instruments

June 26, 2008

Edgar J. Berdahl, Günter Niemeyer, Julius O. Smith III

Contents

1	Introduction	4
1.1	Overview	4
1.2	Musical Acoustics Review	5
1.3	Problem Description	5
1.4	Challenges	6
1.5	Prior Examples	7
1.6	Contributions	7
1.7	Resonance Model Of Instrument	8
1.8	Waveguide Model	8
1.9	Passive Control	9
1.9.1	Definition	9
1.10	Positive Real Functions	9
1.11	Controller Delay	11
1.12	Organization	11
2	Linear Control of Acoustic Instruments	12
2.1	Passive Control	12
2.1.1	Proportional-Integral-Derivative (PID) Control	12
2.1.2	Why changing the resonance frequency is difficult	16
2.1.3	Controller Phase Response	17
2.1.4	Bandpass Control	17
2.1.5	Notch Filter Control	19
2.1.6	Alternating Filter	20
2.1.7	Teleoperator-Based Control	26
2.1.8	Passive Multiple-Input Multiple-Output Controllers	27
2.1.9	Noncollocation Effects	33
2.2	Nonpassive Controllers Motivated By Passive Controllers	41
2.2.1	Negative Spring Constant	41
2.2.2	Negative Damping	42
2.2.3	Negative Mass	42
2.2.4	Integral Control	42
2.3	Higher-Order Control Methods	42
2.3.1	Non-Collocated Traveling Wave-Based Control	42
2.3.2	Collocated Traveling Wave-Based Control	49
2.3.3	Termination-Based Control	50
2.3.4	Loop gain limitations in the laboratory	54
2.3.5	Active Termination	58

3	Self-Sustaining Control of Acoustic Instruments	60
3.1	Introduction	60
3.1.1	Self-Sustaining Oscillators	60
3.1.2	Sources of strong nonlinearity in acoustic musical instruments are often localized . . .	61
3.2	Nonlinear PID Control	61
3.2.1	Damper	62
3.2.2	Spring	62
3.2.3	Lyapunov Stability Of Origin	64
3.2.4	Implications Of Passivity	64
3.3	Application to Musical Self-Sustaining Oscillators	65
3.4	Event-Based Control	67
3.4.1	Approach	67
3.4.2	Analysis	68
3.4.3	Event Detector	72
3.5	RMS Level-Tracking Controllers	73
3.5.1	Dynamic Range Limiter Control	73
3.5.2	RMS Level-Tracking Controllers	74
3.5.3	Convergence	75
3.5.4	Harmonic Distortion	77
3.6	Adaptive RMS Level-Tracking Controllers	79
3.6.1	Multiple Equilibria Example	80
3.6.2	Multiple Band-Pass Filter Effect Example	80
3.6.3	Resonant Ring Modulation Example	81
3.7	Summary	83
4	Conclusions	83
A	Toolbox for the Feedback Control of Sound	85
A.1	Introduction	85
A.2	Open Source Solution	86
A.3	Operation	87
A.4	Estimated PDFs Revealing Jitter Performance	90
A.5	Conclusion	92
A.6	Initialization Files	93
A.6.1	grub.conf	93
A.6.2	/usr/realtime/init_rtai	93
A.6.3	/usr/realtime/start_rtai	93
A.6.4	/usr/local/comedi/load_cmods	93
A.6.5	kcomedilib.c missing functionality	94

B	Actuator and Sensor Design	94
B.1	Harmonic Distortion	94
B.2	Actuators	95
B.2.1	Lorentz' Law Type	95
B.2.2	Reluctance Force Type	96
B.3	Sensors	99
B.3.1	Optical Displacement Sensors	99
B.3.2	Piezoelectric Sensors	102
B.3.3	Sensor THD Estimation	104
B.4	Collocation Considerations	105
B.5	Active String Termination Concept	106
C	Musical Instruments	106
C.1	Electromagnetically-Prepared Piano	106
C.2	Monochord	110
C.3	One String Guitar	111
C.4	Feedback Guitar	112

1 Introduction

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 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. Perhaps interface constancy combined with a wide array of possible instrument sounds presents the strongest argument for the utility of a feedback controlled acoustic musical instrument.

In this document, 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 [52]. Hence, we do not specifically study the discretization problem in this thesis. Readers interested in the implementation details may consult the source code for many of the digital controllers (see Appendix A).

Many controllers discussed are directly applicable to any dissipative acoustic musical instruments. However, our laboratory experiments have been performed in the context of a vibrating string, so some derivations are particular to one-dimensional waveguides or even vibrating strings.




Figure 1: Musician interacting with a feedback controlled acoustic musical instrument

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 [19]. 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.

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 significant detail [26]. 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 3. Consequently, the usual modal analysis techniques are useful for studying the modal shapes and eigenvalues [26]. 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 2), and 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.

Most acoustic musical instruments are harmonic or nearly harmonic. That is, the significant frequencies of vibration are a subset of the harmonic series $f_0, 2f_0, 3f_0, \dots$, where the fundamental frequency f_0 is the largest possible frequency explaining the harmonic modal response. The *pitch* is defined to be what a listener perceives the fundamental frequency to be. 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 the fundamental frequency f_0 . In more complicated situations, the concept of pitch may even be ambiguous: different listeners will perceive significantly different pitches [25].

1.3 Problem Description

Figure 1 depicts how a musician interacts with a feedback controlled acoustic musical instrument [14]. 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,




Figure 2: Simple block diagram for feedback control of an acoustic musical instrument

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 pitch.

The lower part of the block diagram 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 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 shows a simplified version of Figure 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 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.

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. Furthermore, sensing the fundamental frequency of an acoustic musical instrument can be nontrivial [57].
- Control engineers often use 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 [27].
- The modes for acoustic musical instruments typically have long decay times (see Table 2), 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 noise and nonlinearity (see Appendix B).

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 [36] and Sustainiac [53] have mitigated this problem using electromagnetic actuators. Similarly, Weinreich and Caussé have electromagnetically induced the Helmholtz “stick-slip” bowing motion in a vibrating string [72].

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 [30]. Hanagud boosted the second plate mode in an inexpensive acoustic guitar to make it sound more like an expensive acoustic guitar [32]. Chollet, Aeberli, and Charles Besnainou actively tuned the Helmholtz body resonant frequency of a guitar [20].

Besnainou also applied active control to a violin, a snare drum, a pipe organ, and a marimba bar [14]. Boutin has changed the damping time, fundamental frequency, and gain of a xylophone bar using Proportional-Integral-Derivative (PID) control [16][17]. Boutin has also applied PID control to control the physical behavior of a violin bridge [18]. Other researchers have actively controlled drums as well [42][71].

1.6 Contributions

In Chapter 2, we use the positive real formulation of passivity to show why the traditional PID control approach works well in practice. Then we derive controllers using the positive real framework to extend the set of controllers available for feedback control of acoustic musical instruments. These controllers include the bandpass, notch, alternating filter, gyrator, and Multiple Input Multiple Output PID controllers. We briefly explain a few practical nonpassive controllers that are motivated by positive real controllers. We formulate the problem statement of teleoperator-based control of acoustic musical instruments. Next, we analyze noncollocation effects in the context of sensors and actuators that we use in our laboratory. We derive an analytical expression revealing how nonidealities in the sensors and actuators prevent laboratory realizations of controllers from being positive real at high frequencies.

Later in Chapter 2, we consider higher-order control methods. We argue that controlling the fundamental frequency of a one-dimensional waveguide using traveling wave-based control is not robust if the sensors are nonlinear and the traveling wave estimator is linear. In particular, controllers attempting to cause long decay times are likely to be especially problematic. Then we consider collocated traveling-wave based feedback control. We show that implementing a non-dispersive controller would require closing both analog and digital feedback loops simultaneously, rendering this control approach less practical in real-world applications. Next we derive a class of simple controllers for controlling the fundamental frequency of a waveguide. As a special case, we explain how a controller implementing a damper can change the fundamental frequency of a one-dimensional waveguide in theory. Finally we show that the loop gains required to change the fundamental frequency for musical applications with linear controllers are so large that we will not be able to achieve them in the laboratory given our current equipment. Finally, we consider the specific case of creating an active vibrating string termination, and we reduce the problem to that of traveling wave-based control. Therewith we show that the active termination cannot be robust for plucked string implementations with long decay times given even small degrees of sensor nonlinearity.

In Chapter 3, we investigate nonlinear controllers, most of which are designed for sustaining vibrations indefinitely. The first controllers we consider are motivated by strong nonlinearities present in some traditional acoustic musical instruments. Nonlinear PID control may be constrained to be passive, in which case a Lyapunov stability proof shows that the zero-energy state is a locally asymptotically stable equilibrium point. Event-based control for feedback control of acoustic musical instruments is explained, and its behavior is analyzed in an ideal framework. An example event detector is introduced for use in event-based control. Finally, RMS level-tracking control is derived for sustaining vibrations at a specific RMS level when an audio




Figure 3: Damped harmonic oscillator with mass m , damping parameter R , spring constant k , and external force $F(t)$

effect is placed in the feedback loop. Dynamic range limiter control is shown to be a special case of RMS level-tracking control. The harmonic distortion of the control signal given a second-order plant is compared to that of a Van der Pol oscillator's damping signal. It is shown that the speed of the RMS level detector represents a trade off between harmonic distortion and speed. An adaptive RMS level-tracking controller is presented and demonstrated in the context of example simulations with a virtual vibrating string.

1.7 Resonance Model Of Instrument

In the following chapters, we will 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 parameters R , and a spring with parameter k sharing the same velocity form a damped harmonic oscillator as shown in Figure 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:

$$m\ddot{x} + R\dot{x} + kx = F = -u \quad (1)$$

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

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

See Table 2 for some sample values. In this case, the oscillator is considered lightly damped, and the fundamental frequency

$$f_0 \approx \frac{1}{2\pi} \sqrt{\frac{k}{m}}. \quad (3)$$

1.8 Waveguide Model

In other cases, it is important to consider all of the instrument modes. In this case, we apply a digital waveguide model, which models vibrations of a nondispersive one-dimensional waveguide with inverting reflections at the terminations [39]. For most controllers, we have one sensor and one actuator, so this is the case we consider here in Figure 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 [39]. 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 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.

Now we make the model specific to a vibrating string, although it could be easily adapted to many other one-dimensional waveguide instruments [39]. The displacement $x[n]$ measured by the sensor (see Figure 4)




Figure 4: Digital waveguide model of a vibrating string with one sensor and one actuator




Figure 5: Relative dimensions of string model

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 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,

$$\Delta v_{act}[n] = \frac{F[n]}{2R_0}, \quad (4)$$

where R_0 is the string's wave impedance [39].

1.9 Passive Control

1.9.1 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 [59]. 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 [22]. Practical considerations always prevent controllers from being completely passive, so the analysis of closed loop behavior under extreme conditions becomes more complicated (see Section 1.11 and [22]). 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 instrument under positive real control.

1.10 Positive Real Functions

Before defining positive real functions, we first introduce the Laplace transform. Let $\mathcal{L}\{g\} = G$ be shorthand for the following definition of the unilateral Laplace transform [27]:

$$\mathcal{L}\{g(t)\} = G(s) = \int_{0^-}^{\infty} g(t)e^{-st} dt. \quad (5)$$




Figure 6: Driving point of a system

$g(t)$ is the impulse response of the acoustic musical instrument, and $G(s)$ is the Laplace domain representation of the impulse response [27]. Often the complex variable s is decomposed into real and imaginary parts $s = j\omega + \sigma$.

Figure 6 depicts the idea 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 [69]

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

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



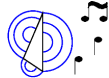
1. $L(s + \epsilon)$ is positive real for all real $\epsilon > 0$.

Now we note some of the properties of positive and strictly positive real functions [69][39]¹:

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. 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 $\int_0^\infty f(t)v(t)dt \geq 0$.

¹Note: The bilinear transform preserves s-domain and z-domain sense positive realness [39].

Table 1: Classification of acoustic musical instrument types

	Stringed instruments
	One-dimensional waveguide instruments
	Dissipative acoustic musical instruments

1.11 Controller Delay

The results from the previous section are quite impressive; however, they are only theoretical. In practice there are several typical sources of nonideality. Besides noncollocation effects (see Section 2.1.9), sensor noise, sensor and actuator nonlinearity, inductive coupling between sensors and actuators (see Appendix B), controller delay can be a significant impediment to desirable closed-loop performance.

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 A).

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 $-\infty$ even if the desired $K(s)$ is positive real.

$$\lim_{\omega \rightarrow \infty} \angle(K(j\omega)e^{-j\omega D}) = -\infty \quad (6)$$

1.12 Organization

In the following chapters, 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 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 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 1.

Figure 7 provides an overview of the controller classes covered in this document. 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.


As discussed in Chapter 3, passive nonlinear controllers are guaranteed to be stable when applied to dissipative musical instruments (see Figure 7). 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.

2 Linear Control of Acoustic 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 passive linear controllers, and we explain how sensor and actuator non-collocation affects can make controllers nonpassive at high frequencies. Then we move on to discuss a few lower-order nonpassive controllers as well as higher-order controllers and their limitations.

2.1 Passive Control



Passive linear controllers satisfy the positive real property (see Section 1.10). 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 8 of the positive real controllers that we will investigate in Section 2.1.

2.1.1 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) [27]. If we measure the velocity, the actuator is collocated with the sensor, and the instrument consists of a single resonance, 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 (7) 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}):

$$u \triangleq P_{DD}\ddot{x} + P_D\dot{x} + P_Px. \quad (7)$$

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

$$K_{PID}(s) = \frac{U(s)}{V(s)} = P_{DD}s + P_D + P_P/s. \quad (8)$$

Substituting (7) into the dynamics (1), we obtain a differential equation describing the closed-loop system [14].

$$(m + P_{DD})\ddot{x} + (R + P_D)\dot{x} + (k + P_P)x = 0 \quad (9)$$

When the new equivalent mass, damping parameter, and spring constant are substituted into (2) and (3), we obtain the new decay time constant $\hat{\tau}$ and the new resonance frequency \hat{f}_0 :

$$\hat{\tau} = \frac{2(m + P_{DD})}{R + P_D} \quad (10)$$

$$\hat{f}_0 \approx \frac{1}{2\pi} \sqrt{\frac{k + P_P}{m + P_{DD}}} \quad (11)$$





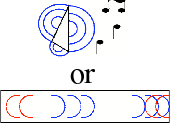

Class	Analysis Class	Linear?	Self-Sustaining?	Passive?	Mechanical Analog	Useful For	Pros	Cons	Implemented?	Simulated?
Positive Real		Y	N	Y	Exists	Controlling musical instruments when details about the dynamics of the instrument (e.g. fundamental frequency) are unknown.	In theory, guaranteed to be stable 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 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 to estimate traveling wave states.	N	Y
Termination-Based		Y	N	N	Does not exist in general	Controlling the fundamental frequency of 1-D waveguide instruments.	Only one sensor is required.	Relatively large loop gains are required in practical contexts.	N	Y
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	 or 	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 7: Overview of controller classes covered in this document

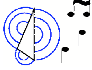
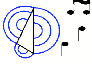


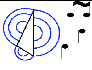

Name	Analysis Class	Linear?	Self-Sustaining?	Passive?	Mechanical Analog	Useful For	Pros	Cons	Implemented?	Simulated?
Proportional-Integral-Derivative (PID) Control		Y	N	Y	Damped harmonic resonator consisting of a mass, spring, and damper	Changing the overall damping; altering the damping and resonance 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 8: Summary of positive real controllers covered in Section 2.1




Figure 9: Collocated control of a vibrating string with a collocated sensor (drawn as microphone) and an actuator (drawn as loudspeaker) (top); Mechanical analog of PID control (bottom)

Differentiating the displacement twice to obtain an estimate of the acceleration is often too noisy in practice [27], 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 9 (top), where $K(s)$ represents the controller. The microphone symbol represents the sensor, and the loudspeaker symbol represents the actuator. Figure 9 (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 [69].


In the laboratory, we have verified that even given a more complicated instrument with multiple resonances, P_D still adjusts the damping, and P_P adjusts the resonance frequencies of the modes.

Given the positive real interpretation from Section 1.10, we know that velocity and force are in some sense the fundamental variables describing 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 causes damping of all of the modes [65].

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 9 (bottom). We will need to search further in order to find a controller capable of changing 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 [9]. 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 A).

2.1.2 Why changing the resonance frequency is difficult



Consider a resonance with a long decay time constant relative to its period. Changing the resonance frequency is more difficult than changing the damping. In particular, a control signal Q larger in amplitude is required to change the damping rate in comparison with changing the resonance frequency squared, where Q is the quality factor of the uncontrolled resonance.

The resonance here is modeled with the mass-spring-damper model from Section 1.7. We assume that the dynamics of the open-loop system are changed only slightly by control, so we can assume that the closed-loop system is still underdamped. Over a small enough time-scale, we can take the displacement $x(t)$ to be a sinusoidal oscillation with amplitude A :

$$x(t) = A\cos(\omega t). \quad (12)$$

We define $\langle \cdot \rangle$ to be the RMS level operator.

$$\langle x(t) \rangle = \sqrt{\int_0^{2\pi/\omega} x(t)^2 dt} = \frac{A}{\sqrt{2}} \quad (13)$$

It follows that $\langle \dot{x} \rangle = \frac{A\omega}{\sqrt{2}}$. The damping rate is the inverse of the time constant τ [26]. Consider the quantity $\frac{\partial \frac{1}{\tau}}{\partial P_D}$ describing how much the damping rate changes as P_D changes while $P_P = 0$ and $P_{DD} = 0$. This quantity makes the most sense when compared with the uncontrolled damping rate, so we form the following ratio:

$$\frac{\frac{\partial \frac{1}{\tau}}{\partial P_D}}{\frac{1}{\tau}|_{P_D=0}} = \frac{\frac{\partial \frac{1}{\tau}}{\partial P_D}}{\frac{1}{\tau}} = -\frac{1}{R} \quad (14)$$

Next we form the corresponding ratio for the change in the resonance frequency squared \hat{f}_0^2 .

$$\frac{\frac{\partial \hat{f}_0^2}{\partial P_P}}{\hat{f}_0^2|_{P_P=0}} = \frac{\frac{\partial \hat{f}_0^2}{\partial P_P}}{\hat{f}_0^2} = -\frac{1}{k} \quad (15)$$

Since displacement feedback changes the resonance frequency, $\langle x \rangle / (\frac{\partial \hat{f}_0^2}{\partial P_P} / \hat{f}_0^2)$ represents the cost of changing the resonance frequency squared. Similarly, since velocity feedback changes the damping rate, the expression $\langle \dot{x} \rangle / (\frac{\partial \frac{1}{\tau}}{\partial P_D} / \frac{1}{\tau})$ represents the cost of changing the damping. Finally, to compare the cost of changing the resonance frequency squared with that of changing the damping rate, we form a ratio.

$$\frac{\langle x \rangle / (\frac{\partial \hat{f}_0^2}{\partial P_P} / \hat{f}_0^2)}{\langle \dot{x} \rangle / (\frac{\partial \frac{1}{\tau}}{\partial P_D} / \frac{1}{\tau})} = \frac{k}{R\omega} = \frac{\omega m}{R} \triangleq Q \quad (16)$$


(16) implies that changing the resonance frequency squared is Q times more difficult than changing the damping rate in the sense that

1. The amplitude of the control signal is Q times larger.
2. The power amplifier for the actuator must be able to produce output currents that are Q times larger.
3. The control power needed is Q^2 times larger.

Musical resonating modes generally have large Q . Whenever $Q \gg \frac{1}{2}$, the decay time is approximately Q periods [38]. Since most musical resonating modes oscillate over very many time periods,² $Q \gg 1$. Hence, the Q in (16) is usually large.

²See Table 2 for more intuition.

2.1.3 Controller Phase Response




From the behavior of the previous controllers, we can already draw some conclusions. 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 $\angle K_{PID}(j\omega) = -\pi/2$ or $\omega > 0$.
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 2.2.1. 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, band pass filters, notch filters, feedforward comb filters, and alternating filters. Many of these controllers are discussed in the next few sections.

2.1.4 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:

$$K_{bp}(s) = \frac{\frac{\omega_c s}{Q}}{s^2 + \frac{\omega_c s}{Q} + \omega_c^2}. \quad (17)$$

The frequency response of the bandpass filter with $w_c = 2\pi 200\text{Hz}$ and $Q = 20$ is shown in Figure 10. The effect of the controller can be determined chiefly by looking at the phase response (see Figure 10, bottom). Any instrument resonances located very near ω_c will be damped because to $\angle K_{bp}(j\omega_c) = 0$ 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 (17) is shown in Figure 11 in the context of controlling a vibrating string. The existence of the mechanical analog proves that $K_{bp}(s)$ is positive real [69]. 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 11 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.




Figure 10: Frequency response of the bandpass controller given by (17) with $Q = 20$ and $w_c = 2\pi 200\text{Hz}$




Figure 11: Mechanical analog of the bandpass controller given by (17)





Figure 12: Frequency response of the notch filter controller given by (18) with $\omega_c = 2\pi 300\text{Hz}$, $Q = 1$, and $\alpha = 10$

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

$$K_{notch}(s) = \frac{s^2 + \frac{\omega_c s}{\alpha Q} + \omega_c^2}{s^2 + \frac{\omega_c s}{Q} + \omega_c^2}. \quad (18)$$

The frequency response is shown in Figure 12 for $\omega_c = 2\pi 300\text{Hz}$, $Q = 1$, and $\alpha = 10$. Q is no longer the quality factor as defined in (16), 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 13 (right) in the s -plane. Restricting $\alpha > 1$ ensures that the poles lie behind the zeros with respect to the imaginary axis of the s -plane (see Figure 13).

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 (18) in series in the signal processing chain.




Figure 13: Mechanical analog of the notch controller given by (18) (left); Pole-zero plot of (18) (right)

2.1.6 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, the alternating filter should change the pitch while still mostly preserving the harmonicity of the waveguide. In light of the design criteria in Section 2.1.3, 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 in Figure 14 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 14, bottom).

Since the controller is complicated, we design it directly in the z -domain rather than first in the s -domain and then applying the bilinear transformation [39]. The definition of z -domain positive realness is similar to the s -domain definition given in Section 1.10 [39].

To create the response shown in Figure 14, the controller needs to consist predominantly of complex pole and zero pairs placed very near the unit circle as shown in Figure 15. The poles and zeros need to be interlaced in order to satisfy the positive real property [39]. 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 15. 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 16).

We tested the controller on a digital waveguide simulation of a vibrating string. 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 A). A zoomed-out version of $K_{af}(z)$'s frequency response is shown in Figure 17. The dashed black lines in Figure 17 (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 18 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 18, 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 19). 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




Figure 14: Frequency response of $K_{af}(z)$ where the x 's mark where the transfer function overlaps with the resonances of the instrument with $f_0 = 100\text{Hz}$




Figure 15: Pole zero plot of the digital implementation of $K_{af}(z)$ (zoomed in)




Figure 16: Pole zero plot of the digital implementation of $K_{af}(z)$




Figure 17: Frequency response of the alternating filter $K_{af}(z)$




Figure 18: Frequency response $K_{af}(z)z^{-1}$ incorporating the controller delay





Figure 19: Fourier transform of windowed open loop response (solid blue) and closed loop response (dash-dotted red)

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 presented here is positive real, and it is less sensitive to many of the details of the instrument model. In fact, in simulation we found the controller still functioned effectively given sensor and actuator nonlinearity on the order of -20dB THD.

2.1.7 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 [19]. These considerations motivate the development of teleoperator-based control for the feedback control of acoustic musical instruments.

Figure 20 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 instrument is linked to the driving point of a virtual dissipative musical instrument by way of the spring k and damper R . The spring binds the physical and virtual elements together, while the damper provides additional damping, which is often helpful in ensuring stability in practice.

The controller depicted in Figure 20 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




Figure 20: Teleoperator-based control of a vibrating string

connects the master robot to a slave robot, which typically completes a task requiring the human’s assistance [33].

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.

Note that the bandpass controller from Section 2.1.4 is analogous to a lightly-damped harmonic oscillator centered about ω_c . However, in contrast with teleoperator-based control, when the simple bandpass controller is applied directly, the bandpass controller somewhat counterintuitively damps resonances near ω_c instead of inducing a closed-loop resonance at ω_c .

2.1.8 Passive 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, such as stiffness elements. We noted that increasing the stiffness of the string locally effectively increased the frequency of the lowest partial (see Section 2.1.1). This is physically different than increasing the tension of the entire string in a distributed sense, which directly affects the fundamental frequency [26].

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 passive types controllers for collocated multiple-input multiple-output (MIMO) geometries.

Many passive controllers can be specified by using the positive real formulation, which we now generalize 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 [24]:

1. All elements of $K(s)$ are analytic for $Re\{s\} > 0$, 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 $Re\{s\} > 0$, which is equivalent to stating that $K(s)$ is Hermitian positive semidefinite for $Re\{s\} > 0$.

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 21: Mechanical analog for the controller described by (19)




Figure 22: The 2-port gyrator element

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 apply any of the previously discussed passive 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 21 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

$$K_{PD}(s) = \frac{P_D s + P_P}{s} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \quad (19)$$

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 passive elements such as transformers and gyrators. The 2-port transformer is not likely to be of use in the feedback control of musical instruments because it merely scales gains, so we do not consider the transformer further here. On the other hand, the 2-port gyrator, as shown in Figure 22, has many useful properties [69]. 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 system 1 $P_{added,1}$ is the same as the power removed from system 2 $P_{removed,2}$.

$$P_{added,1} = F_1 V_1 = \frac{V_2}{u} (-u F_2) = -F_2 V_2 = P_{removed,2} \quad (20)$$




Figure 23: Block diagram for gyrator connecting two independent second-order systems

Equivalently, passivity of the gyrator can be shown starting from

$$K_{gyrator}(s) = \begin{bmatrix} 0 & u \\ -u & 0 \end{bmatrix} \quad (21)$$

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 systems are symmetric because they do not include gyrators [55].

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 [39], but for simplicity, we will assume that the coupling is zero. We model each axis as a simple second-order system (see Section 2.1.1). This model is really an oversimplification, but it is not unreasonable for initial analysis. We excite the first system with an additional force $F_{in}(s)$:

$$\frac{m_1 s^2 + R_1 s + K_1}{s} V_1(s) = u V_2(s) + F_{in}(s) \quad (22)$$

$$\frac{m_2 s^2 + R_2 s + K_2}{s} V_2(s) = -u V_1(s). \quad (23)$$

We arrive at the standard feedback control block diagram shown in Figure 23.

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 24. 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 24. In 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 [39].

GYRATOR BETWEEN TWO 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 (21)




Figure 24: Root locus for gyration applied to two independent lightly-damped harmonic oscillators




Figure 25: Gyrotor connected between the points A and B along a vibrating string

be applied between the points A and B on a vibrating string as shown in 25. The reflection coefficient for velocity traveling waves impinging on the physically almost completely rigid terminations is assumed to be $-p$. Figure 25 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 $q = \frac{u}{2R_0}$ to simplify the following expressions.

$$\frac{V_A^{+R}(f)}{V_A^{-R}(f)} = \frac{-pe^{-j4\pi f\tau_A} + qe^{j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}{1 - qe^{-j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})} \quad (24)$$

$$\frac{V_B^{+R}(f)}{V_B^{-R}(f)} = \frac{\frac{-pe^{-j4\pi f\tau_A} + qe^{j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}{1 - qe^{-j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})} (e^{-j2\pi f\tau_B} - q) - q}{e^{j2\pi f\tau_B} + q + q \frac{-pe^{-j4\pi f\tau_A} + qe^{j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}{1 - qe^{-j2\pi f\tau_B}(1 - pe^{-j4\pi f\tau_A})}} \quad (25)$$

Taking the limit of (24) and (25), we obtain the following for the case of infinite control gain:

$$\lim_{|u| \rightarrow \infty} \frac{V_A^{+R}(f)}{V_A^{-R}(f)} = \lim_{|u| \rightarrow \infty} \frac{V_B^{+R}(f)}{V_B^{-R}(f)} = -1. \quad (26)$$

Furthermore by symmetry, we also have the following, which in conjunction with (26) imply that the system becomes three isolated vibrating string segments at infinite control gain:

$$\lim_{|u| \rightarrow \infty} \frac{V_A^{-L}(f)}{V_A^{+L}(f)} = \lim_{|u| \rightarrow \infty} \frac{V_B^{-L}(f)}{V_B^{+L}(f)} = -1. \quad (27)$$

Finally, let the mobility of the controlled string at junction B be $Y_B(f)$.

$$Y_B(f) = \frac{1}{2R_0} \left(\frac{1 - pe^{-j4\pi f\tau_C}}{1 + p \left(\frac{V_B^{+R}(f)}{V_B^{-R}(f)} \right) e^{-j4\pi f\tau_C}} \right) \left(1 + \frac{V_B^{+R}(f)}{V_B^{-R}(f)} \right). \quad (28)$$

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 $\frac{1}{2(\tau_A + \tau_B + \tau_C)} = 122\text{Hz}$. $p = 0.9$ is chosen artificially small so that the plot is easier to inspect visually; however, it does not alter the qualitative results: the gyrotor causes the resonance frequencies of some of the modes to increase and of other modes to decrease (see Figure 26). The decay times of some of the modes increase, while the decay times of other modes decrease. Thus the effect is interesting, but perhaps not so useful in practice as it causes inharmonicity of the string. We have verified the behavior qualitatively in simulation using the digital waveguide model introduced in Section 1.8. Note that since the closed-loop system remains passive, the mobility $Y_B(f)$ must remain positive real, so the phase response lies within $[-\pi/2 \ \pi/2]$ as in Figure 26 (bottom).





Figure 26: Mobility at junction B of the system shown in Figure 25




Figure 27: Actuator (blue), sensor (red), and waveguide (black) orientation where $-L_{S-}/2 \geq -L_A/2$ (case I)

2.1.9 Noncollocation Effects



While perfectly collocating a sensor and an actuator ensures that $G(s)$ is positive real, collocating a sensor and an actuator in practice can be challenging. For instance, placing a sensor and an actuator closely to one another may result in direct inductive coupling (see Section B.4). In addition, if either the sensor or actuator does not operate purely at a point, the system becomes non-positive real as seen through the actuator and sensor. This undesirable side effect results from propagation delays in the structure being controlled, and it greatly affects considerations when designing actuators and sensors for control. In this section, we examine the specific effects of an actuator and sensor both operating upon a finite length of a one-dimensional waveguide.

Consider the sensor and actuator orientation shown in Figure 27. Here we assume that sensor and actuator are interfaced with a one-dimensional lossless waveguide characterized by wave impedance R_0 and wave speed c . For convenience, we will assume that the actuator exerts a the commanded force on the waveguide over the lateral range $x \in [-\frac{L_A}{2} \quad \frac{L_A}{2}]$ (see the blue line segment in Figure 27), where x denotes a position along the waveguide. In contrast, the sensor is not necessarily laterally centered. It measures the average of the velocity over the lateral range $x \in [-\frac{L_{S-}}{2} \quad \frac{L_{S+}}{2}]$ (see Figure 27, in red). Let L_S be the total length of the waveguide over which the sensor measures:

$$L_S = \frac{L_{S+} + L_{S-}}{2} \quad (29)$$

Next we determine the generalized mobility $H_{sa}(f)$, which is the transfer function between the sensed velocity $V_{sens}(f)$ and the actuator force $F(f)$. If the sensor and actuator were perfectly collocated with $L_A \rightarrow 0$ and $L_S \rightarrow 0$, then $H_{sa}(f)$ would represent a true mobility [26]. We assume that the waveguide is infinitely long supporting no reflections, so ideally in the collocated case we would have $H_{sa}(f) = \frac{1}{2R_0}$, allowing $H_{sa}(f)$ to be positive real.

CASE I

A straight-forward way to find $H_{sa}(f)$ for the orientation shown in Figure 27 involves breaking the actuator down into a series of identical actuators (see the “multi-actuator” in Figure 28, blue). Let the i th actuator be modeled approximately as an ideal point actuator inducing rightward and leftward traveling velocity waves [39]:

$$V_{iR} = \frac{F_i}{2R_0} \quad \text{and} \quad V_{iL} = \frac{F_i}{2R_0}. \quad (30)$$

The velocity waves propagate according to the wave speed c , and let each sensing element sense each delayed component (see the “multi-sensor” in Figure 28, red). The final sensed signal $V_{sens}(f)$ is the sum of the velocities measured at each sensor. Finally, take the limit as of the number of elements in the multi-actuator and multi-sensor approach infinity, while they still correspond to the same finite lengths of the waveguide




Figure 28: Multi-actuator (blue), multi-sensor (red), and waveguide (black, middle)

L_A and L_S , respectively. Without loss of generality, let $-L_{S-}/2 \geq -L_A/2$ and $L_{S+}/2 \leq L_A/2$ as shown in Figure 28. If we further scale by the actuator and sensor lengths appropriately, we arrive at an integral expression, where propagation delays in the waveguide are modeled using the Fourier transform of a delayed impulse $\delta(t - \tau) \longleftrightarrow e^{-j2\pi f\tau}$ [54]:

$$V_{sens}(f) = \frac{1}{2R_0} \int_{w=-\frac{L_{S-}}{2}}^{\frac{L_{S+}}{2}} \int_{u=-\frac{L_A}{2}}^{\frac{L_A}{2}} \frac{1}{L_A L_S} F(f) e^{-j2\pi f \frac{|w-u|}{c}} du dw. \quad (31)$$

$$V_{sens}(f) = \frac{F(f)}{2R_0 L_A L_S} \int_{w=-\frac{L_{S-}}{2}}^{\frac{L_{S+}}{2}} \left[\int_{u=-\frac{L_A}{2}}^w e^{-j2\pi f \frac{w-u}{c}} du + \int_{u=w}^{\frac{L_A}{2}} e^{-j2\pi f \frac{u-w}{c}} du \right] dw \quad (32)$$

After extensive manipulation, we arrive at the following generalized mobility of the sensor, waveguide, actuator system:

$$H_{sa}(f) = \frac{V_{sens}(f)}{F(f)} = \frac{L_{S+} + L_{S-} - e^{j\pi f L_A/c} (L_{S+} \text{sinc}(f L_{S+}/c) + L_{S-} \text{sinc}(f L_{S-}/c))}{2R_0 L_A L_S (j2\pi f/c)}. \quad (33)$$

The expression for the transfer function becomes simpler if we let $L_{S+} = L_{S-} = L_S$.

$$H_{sa,sym}(f) = \frac{1 - e^{j\pi f L_A/c} \text{sinc}(f L_S/c)}{R_0 L_A (j2\pi f/c)} \quad (34)$$

Figure 29 shows the symmetric expression given in (34) normalized by the wave impedance R_0 and evaluated for $c = 111\text{m/s}$ and $L_S = 4\text{mm}$. Figure 29 shows that $L_A = 0.64\text{cm}$ ($0.25''$) will allow larger system bandwidths than $L_A = 1.9\text{cm}$ ($0.75''$), suggesting that we should consider using the smaller pair of magnets in our actuator rather than the larger pair (see Section B.4). In general, making L_A and L_S smaller increases the bandwidth.

Since the phase response is linear at lower frequencies, we develop a parsimonious model $H_{delay}(f) = e^{-j2\pi f\tau}$ such that $\angle H_{delay}(f) \approx \angle H_{sa,sym}(f)$ for small f . Note that the model is perhaps overly parsimonious in that the magnitude is taken to be constant; however, it is still interesting to consider the approximate length of the noncollocation delay.

$$\tau = \lim_{f \rightarrow 0} \angle H_{sa,sym}(f) = -\frac{\pi}{2} + \lim_{f \rightarrow 0} \angle(1 - e^{j\pi f L_A/c} \text{sinc}(f \frac{L_S}{c})). \quad (35)$$

$$\tau = -\frac{\pi}{2} + \lim_{f \rightarrow 0} \arctan \frac{\sin(\pi f \frac{L_A}{c}) \text{sinc}(f \frac{L_S}{c})}{1 - \cos(\pi f \frac{L_A}{c}) \text{sinc}(f \frac{L_S}{c})} \quad (36)$$

If we then assume that $L_S \ll L_A$, then we can simplify further.

$$\tau = -\frac{\pi}{2} + \lim_{f \rightarrow 0} \arctan \frac{\sin(\pi f \frac{L_A}{c})}{1 - \cos(\pi f \frac{L_A}{c})} \quad (37)$$




Figure 29: $R_0 H_{sa,sym}(f)$ plotted for $L_S = 4\text{mm}$ given $L_A = 1.9\text{cm}$ (0.75", blue) and $L_A = 0.64\text{cm}$ (0.25", red)




Figure 30: Comparison of the calculated and high-frequency approximate formulas for $R_0 H_{sa,sym}(f)$ for $L_A = 1.9\text{cm}$ ($0.75''$) and $L_S = 4\text{mm}$

$$\tau = -\frac{\pi}{2} + \lim_{f \rightarrow 0} \left(\frac{\pi}{2} - \pi f \frac{L_A}{2c} \right) = -2\pi f \frac{L_A}{4c} \quad (38)$$

$$H_{delay}(f) = e^{-j2\pi \frac{L_A}{4c} f} \quad (39)$$

The model given by (39) implies that when the actuator is much wider than the sensor, and when the center of the actuator and sensor are aligned, the noncollocation delay is roughly equal to the propagation delay along $\frac{1}{4}$ of the actuator window.

We know that the propagation delays should be negligible at DC, so the sensor and actuator behave ideally at very low frequencies:

$$\lim_{f \rightarrow 0} H_{sa}(f) = \frac{1}{2R_0}. \quad (40)$$

At high frequencies, the inertia of the waveguide dominates, so then $H_{sa}(f)$ behaves like an integrator:

$$\lim_{f \rightarrow \infty} H_{sa}(f) = \frac{c}{j2R_0\pi f L_A}. \quad (41)$$

For the special case where $L_A = 0.75''$ (1.9cm) and $L_S = 4\text{mm}$, the calculated and high-frequency approximate formulas for $H_{sa,sym}(f)$ are shown in Figure 30. As expected, the high-frequency approximation becomes more accurate at higher frequencies.




Figure 31: Actuator (blue), sensor (red), and waveguide (black) orientation, where $\frac{L_{S+}}{2} > \frac{L_A}{2}$ and $\frac{-L_{S-}}{2} \leq \frac{L_A}{2}$ (**case II**)

CASE II

For completeness, we also consider the orientation shown in Figure 31, where $\frac{L_{S+}}{2} > \frac{L_A}{2}$ but we still have $\frac{-L_{S-}}{2} \leq \frac{L_A}{2}$. Because of the new boundary conditions, the integral given in (42) becomes more complicated.

$$V_{sens,II}(f) = \frac{F(f)}{2R_0L_AL_S} \int_{w=-\frac{L_{S-}}{2}}^{\frac{L_A}{2}} \left[\int_{u=-\frac{L_A}{2}}^w e^{-j2\pi f \frac{w-u}{c}} du + \int_{u=w}^{\frac{L_A}{2}} e^{-j2\pi f \frac{u-w}{c}} du \right] dw + \int_{w=\frac{L_A}{2}}^{\frac{L_{S+}}{2}} \int_{u=-\frac{L_A}{2}}^{\frac{L_A}{2}} e^{-j2\pi f \frac{w-u}{c}} dudw \quad (42)$$

$$H_{sa,II}(f) = \frac{L_A(1 - e^{-j\pi f L_{S+}/c} \text{sinc}(f \frac{L_A}{c})) + L_{S-}(1 - e^{-j\pi f \frac{L_A}{c}} \text{sinc}(f \frac{L_{S-}}{c}))}{2R_0L_AL_S(j2\pi f/c)} \quad (43)$$

CASE III

In the final case, $-\frac{L_{S-}}{2} \geq \frac{L_A}{2}$, so the sensor does not overlapping with the actuator at all (see Figure 31). In other words, the sensor and actuator are completely non-located. The same integral approach may be applied to find $H_{sa,III}(f)$, but there is a simpler approach. Without loss of generality, the only wave that matters is the rightward traveling wave. The average propagation delay between the actuator and the sensor

$$\tau_{pd} = \frac{\frac{L_{S+}}{2} - \frac{L_{S-}}{2}}{2c} = \frac{L_{S+} - L_{S-}}{4c}. \quad (44)$$

Outside of the actuator window, the rightward traveling wave $v^+(t)$ arriving at the center of the sensor is the time-delayed, scaled input force signal convolved by rectangular actuator windowing function [39][62]:

$$v^+(t) = \frac{f(t - \tau_{pd})}{2R_0L_AL_S} * \Pi\left(\frac{t}{L_A/c}\right), \quad (45)$$

where

$$\Pi(t) = \begin{cases} 1 & \text{if } |t| \leq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases} \quad (46)$$

Next we take the Fourier transform of both sides of (45), taking use of the well-known transform $\Pi(\frac{t}{L_A/c}) \longleftrightarrow \text{sinc}(f \frac{L_A}{c})$ [54].

$$V^+(f) = \frac{F(f)e^{-j2\pi f \tau_{pd}}}{2R_0L_AL_S} \text{sinc}(f \frac{L_A}{c}) \quad (47)$$




Figure 32: Actuator (blue), sensor (red), and waveguide (black) orientation, where $-\frac{L_{S-}}{2} \geq \frac{L_A}{2}$ (**case III**)

Finally, the sensed signal is $V^+(f)$ filtered according to the sensor windowing function $\prod(\frac{t}{L_S/c}) \longleftrightarrow \text{sinc}(f\frac{L_S}{c})$.

$$V_{sens,III}(f) = V^+(f)\text{sinc}(f\frac{L_S}{c}) = \frac{F(f)e^{-j2\pi f\tau_{pd}}}{2R_0L_AL_S} \cdot \text{sinc}(f\frac{L_A}{c})\text{sinc}(f\frac{L_S}{c}) \quad (48)$$

$H_{sa,III}(f)$ is a linear-phase filter, so we can actually make use of the known propagation delay to determine the zero-phase version $H_{sa,III,zp}(f)$ of the filter [52].

$$H_{sa,III,zp}(f) = H_{sa,III}(f)e^{j2\pi f\tau_{pd}} = \frac{1}{2R_0L_AL_S} \text{sinc}(f\frac{L_A}{c})\text{sinc}(f\frac{L_S}{c}) \quad (49)$$

The completely non-collocated zero-phase transfer function $H_{sa,III,zp}(f)$ is compared with the symmetric case $H_{sa,I,sym}(f)$ in Figure 33 for $L_S = 4\text{mm}$ and $L_A = 0.75''$ (1.9cm). The zeros of $H_{sa,III,zp}(f)$ manifest themselves clearly at the locations $f = \frac{cn}{L_A}$ and $f = \frac{cn}{L_S}$ for integers $n \neq 0$ (see Figure 33, top in green). These zeros are pushed back off of the frequency axis for $H_{sa,I,sym}(f)$, but their effects can still be seen in Figure 33, top in blue). From the point of view of $H_{sa,III,zp}(f)$, minimizing the maximum of L_S and L_A is desirable for promoting maximum system bandwidth.

Curiously the following relation holds for only the symmetric case where $L_{S+} = L_{S-} = L_S$:

$$\text{Re}\{H_{sa,I,sym}(f)\} = H_{sa,III,zp}(f). \quad (50)$$

In some situations, propagation delays can sometimes be compensated for by moving part of the system to a different point along the waveguide. The existence of $H_{sa,III,zp}(f)$ is indeed important, and it is indeed one of the motivations behind the traveling wave-based controllers described in Section 2.3.1—ideally these controllers have no phase lag even at high frequencies because they see the traveling wave approaching in advance. However, their behavior is predicated on the approaching wave not being changed by noise, nonlinearity, or external disturbances between the estimation of the traveling wave and its arrival. These nonidealities can adversely affect performance in practical situations (see Section 2.3.1).

OPTIMAL PLACEMENT

Clearly in collocated control techniques, a sensor and its accompanying actuator should not be misaligned by any large amount because then propagation delays will reduce the control system phase margin. However, it turns out that having them perfectly centered is not always best. Given our sensor and actuator components in the laboratory, it turns out theoretically that having the sensor displaced slightly off-center reduces the maximum phase lag without significantly increasing the delay at low frequencies.

We plotted the family of transfer functions $H_{sa}(f)$ for $L_{S-} \in [L_S - \frac{3L_A}{4}]$ in blue in Figure 34. As before, $L_A = 0.75''$ (1.9cm) and $L_S = 4\text{mm}$. Compared to the symmetric case (see Figure 29), many of the phase responses contribute a smaller maximum phase lag. In particular, the red curve for $\frac{L_{S-}}{2} = -0.42\text{mm}$ is quite well behaved. By having some of the zeros apparently pulled further away from the frequency axis, this preferred alignment reduces the range over which the phase response varies.




Figure 33: Symmetric configuration $R_0 H_{sa,I,sym}(f)$ (case I, blue) versus the completely non-collocated zero-phase transfer function $R_0 H_{sa,III,zp}(f)$ (case III, green)




Figure 34: $R_0 H_{sa}(f)$ plotted in blue for various L_{S-} with $L_A = 0.75''$ (1.9cm) and $L_S = 4\text{mm}$. The preferred configuration $\frac{L_{S-}}{2} = -0.42\text{mm}$ is shown in red.





Figure 35: Preferred sensor/actuator alignment

We prefer the sensor/actuator alignment depicted in Figure 35 because it minimizes the maximum phase lag. It furthermore behaves somewhat like a leaky-integrator, so it could be used as the lowest-frequency pole to roll off various controllers (see Figure 30 for a less optimal but nevertheless similar example). While its phase response does exceed the $\pm\frac{\pi}{2}$ boundary allowed by positive-real transfer functions, it is nevertheless close, and so we hope that in practice, we are able to approximately implement positive real controllers.


Finally, consider a waveguide with terminations. Then of course the approximate “mobility” of a collocated sensor/actuator pair changes considerably. However, the direct path between the actuator and sensor is still described by $H_{sa}(f)$, while reflections from the terminations as seen by the controller are filtered by $H_{sa,III,zp}(f)$. This underscores the importance of the transfer functions we have investigated in this section. We are certain that in general, reducing the size of the actuator and sensor windows is desirable. Furthermore, to reduce non-collocation effects, the actuator and sensor should be nearly centered. However, in our case, it turned out to be better to place the sensor slightly off center as shown in Figure 35. Since the centered configuration was slightly sub-optimal for this geometry, we conclude that there are likely other geometries for which perfect centered alignment is also suboptimal.

2.2 Nonpassive Controllers Motivated By Passive Controllers




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

2.2.1 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 (7)). 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 (10) and as observed in the laboratory with a vibrating string, $P_P < 0$ causes the pitch of the controlled instrument to decrease.

2.2.2 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 (11), negative damping may be employed to increase the decay time of instrument resonances.

2.2.3 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 (10) and (11).

2.2.4 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. Since we choose $a > 0$ but $a \approx 0$, we essentially integrate a velocity measurement twice to obtain a quantity similar to the integral of displacement:

$$K_{int}(s) = P_I \frac{a}{s(s+a)} \quad (51)$$

To apply damping with integral control, we choose $P_I < 0$ so that for all frequencies sufficiently far enough away from DC, the controller similarly to a low-pass damper. That is, for $\omega \gg a$, $\angle K_{int}(j\omega) \approx 0$. Nevertheless, $K_{int}(s)$ is not passive 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 1.7 with the help of the root locus diagram shown in Figure 36. 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 very negative loop gains P_I . Therefore a should be chosen such that strong enough stable control can be obtained for the application while $a \ll 2\pi f_l$, where f_l is the resonance frequency of the lowest open-loop instrument mode.

Employing a pole perturbation analysis for small $|P_I|$, we obtain a formula for the new decay time constant [12]:

$$\hat{\tau} \approx \frac{2m}{R + P_I/(4\pi^2 f_0^2)}. \quad (52)$$

2.3 Higher-Order Control Methods

2.3.1 Non-Collocated Traveling Wave-Based Control



Waveguide models, as described in Section 1.8, 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 37. 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




Figure 36: Root locus over $P_I < 0$ for $K_{int}(s)$ applied to a lightly damped harmonic oscillator




Figure 37: Representation of a physical waveguide

eliminated the loop filter $H_{lp}(z)$ so that the partials are perfectly harmonic and have the same decay time constant τ [39]³:

$$\tau = -\frac{1}{f_0 \ln(|g|)}, \quad (53)$$

The fundamental frequency $f_0 = f_S/(A + B + C + D)$, where f_S is the sampling rate in Hz. Since the traveling waves cannot be sensed directly, we must estimate them using arrays of sensors [8]. In this section, we consider how sensor nonlinearity prevents this control approach from being robust.

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 [8]. 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 38. 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 [31]. 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 [8].

To truly bring a virtual waveguide instrument to the physical world, it must be possible to change the fundamental frequency by a significant amount. Then N and $B + C$ must differ considerably. Except for perhaps the lowest mode of vibration, the phase of the injected wave modes will diverge significantly from what they would have been in the open-loop case. The large phase shifts imply that the amount of control power required is relatively large. We may as well damp the physical reflection passively so that g becomes smaller. In practice given a vibrating string, it is hard to make g smaller than 0.8 using passive methods.⁴ Consider the waveguide shown in Figure 39, where the physical reflection has already been damped. We will argue that since implementing an artificial termination is not robust, concurrently damping a physical termination and implementing an artificial termination cannot be robust either.

In the following, we will argue that inducing an artificial reflection is not robust given sensor nonlinearity. Since we cannot robustly induce an artificial reflection, and we save little control power in general whether we choose 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 40. Let the variable we are controlling be v , which is velocity in the case of the vibrating string. The coefficient $g_{requested}$ models the loop gain implemented directly by the controller as a standard gain block. $\frac{\partial \hat{v}_R}{\partial v_R}$ models how the linearized loop gain

³If instead one end of the waveguide is inverting and the other non-inverting, for example in the case of a flute, then actually $\tau = -\frac{1}{2f_0 \ln(|g|)}$.

⁴We typically do this in the laboratory by wrapping felt around the end of the string or placing batting underneath the end of the string. It helps to make the impedance of the damping material increase gradually from a small value to a large value (see [41] p. 175 for a transmission line varying-impedance termination example). We have considered further alternatives involving dashpots and viscous liquids, but we have never tested these.




Figure 38: Configuration for non-collocated traveling wave-based control




Figure 39: Configuration after physical reflection has already been damped.




Figure 40: Configuration where sensor nonlinearity is modeled by $\frac{\partial \hat{v}_R}{\partial v_R}$

is skewed by the sensor nonlinearity. To keep the two 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:

$$\frac{\partial \hat{v}_R}{\partial v_R} \approx \frac{\partial \hat{v}}{\partial v}. \quad (54)$$

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 (55). $\hat{v}(x, t)$ is the distorted sensor measurement of the actual velocity $v(x, t)$ at the position x along the waveguide.

$$\hat{v}(x, t) = v(x, t) + \alpha v^3(x, t) \quad (55)$$

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

$$\frac{\partial \hat{v}(x, t)}{\partial v(x, t)} = 1 + 3\alpha v^2(x, t) \quad (56)$$

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

$$v(x, t) = A \cos \omega t, \quad (57)$$

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

$$\hat{v}(x, t) = A \cos \omega t + \alpha A^3 \cos^3 \omega t \quad (58)$$

$$\hat{v}(x, t) = A \cos \omega t + \frac{\alpha}{4} A^3 (3 \cos \omega t + \cos 3\omega t) \quad (59)$$

$$\hat{v}(x, t) = (A + \frac{3}{4} \alpha A^3) \cos \omega t + \frac{\alpha}{4} A^3 \cos 3\omega t \quad (60)$$

From Equation (60), 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 (55), the dependency is much more complicated. We continue by applying the definition of total harmonic distortion (see Section B.1).

$$k = \sqrt{\frac{\frac{1}{2}(\frac{\alpha}{4}A^3)^2}{\frac{1}{2}[(A + \frac{3}{4}\alpha A^3)^2 + (\frac{\alpha}{4}A^3)^2]}} \quad (61)$$

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

$$k = \sqrt{\frac{\frac{1}{2}(\frac{\alpha}{4}A^3)^2}{\frac{1}{2}(A)^2}} = \frac{\alpha}{4} A^2 \quad (62)$$

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.⁵ That is, we require that

$$|g_{ctrl}| < 1. \quad (63)$$

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

$$|g_{ctrl}| = |g_{requested}| \left| \frac{\partial \hat{v}_R}{\partial v_R} \right| < 1 \quad (64)$$

Then (54) and (56) imply the following bound:

$$|g_{requested}| < \frac{1}{|1 + 3\alpha v^2(x, t)|} \quad (65)$$

A sufficient condition for (65) can be developed by substituting the maximum value of $v^2(x, t)$ for $v^2(x, t)$.

$$|g_{requested}| < \frac{1}{1 + 3\alpha \max\{v^2(x, t)\}} \quad (66)$$

Finally, we substitute k into the above inequality by applying (57) and then (62):

$$|g_{requested}| < \frac{1}{1 + 3\alpha A^2} \quad (67)$$

$$|g_{requested}| < \frac{1}{1 + 12k} \quad (68)$$

Table 2: Relation between distortion and maximum safe waveguide loop gain

k	Value of k	Maximum $g_{requested}$	Decay τ	Description
$k_{velocity}$	-26dB	0.62	0.02	
$k_{displacement}$	-33dB	0.79	0.04	
	-40dB	0.89	0.09	pizzicato pluck
	-60dB	0.988	0.8	normal guitar pluck
	-80dB	0.999	8.3	very long guitar pluck

The value of k that we currently achieve in the laboratory is $k_{velocity}$, as given in Table 2. 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 (55),⁶ we can guarantee that the linearized loop gain will never become unstable, as long as we choose $|g_{requested}| < 0.62$.

This is a large price to pay for controller robustness. Choosing $|g_{requested}| = 0.62$ is equivalent to requesting for a system with time constant $\tau = 0.02$ s at $f_0 = 100$ 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 2). 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_{ctrl} must be chosen small to prevent various harmonics from becoming unstable due to sensor nonlinearity; however, once g_{ctrl} 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 2, line 4).

In contrast, if we allow the traveling wave estimator to be nonlinear, then we can solve the problem trivially. 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 main 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 B.3.1), further frustrating any sensor nonlinear characteristic inversion.

One may ask how Guérard successfully implemented a non-collocated traveling wave-based controller [31]. 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. Furthermore, Guérard states that some instability problems remained even at the end of his study [31]. It is possible that microphone sensor nonlinearity was one of the factors detracting from the robustness of the controlled recorder.

⁵Note that we are not considering how energy is transferred between modes due to nonlinear coupling, which could still drive the system unstable.

⁶Take careful note of how $k_{velocity}$ is determined in Section B.3.1. One additional difference is that the argument assumes we are making a distorted velocity measurement (see 55). In contrast, in our laboratory, we are actually making a distorted displacement measurement, which we differentiate 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.




Figure 41: Configuration for collocated traveling wave-based control

2.3.2 Collocated Traveling Wave-Based Control



In collocated traveling wave-based control, we assume that we know the rightward traveling wave at a certain point along the waveguide, and we can actuate the rightward traveling wave at this point as well. An example is shown in Figure 41, where the collocated control point lies in between the A and B traveling wave delay segments.⁷ $C(z)$ is a digital filter implemented by the controller.

Let $C(z)$ be an infinite impulse response (IIR) filter with real coefficients [52]:

$$C(z) = \frac{b_0 + b_1 z^{-1} + \dots + b_M z^{-M}}{1 + a_1 z^{-1} + \dots + a_N z^{-N}}. \quad (69)$$

$$H_{control}(z) = \frac{1}{1 - C(z)} = \frac{1 + a_1 z^{-1} + \dots + a_N z^{-N}}{1 + a_1 z^{-1} + \dots + a_N z^{-N} - b_0 - b_1 z^{-1} - \dots - b_M z^{-M}} \quad (70)$$

If $H_{control}(z)$ is not linear phase, then the higher partials will be detuned [39]. Since only finite impulse response (FIR) filters can be linear phase [52], $H_{control}(z)$ must be an FIR filter to make it possible to avoid detuning the higher partials. Thus, we choose $M = N$ and $b_1 = a_1, b_2 = a_2, \dots, b_N = a_N$.

$$H_{control}(z) = \frac{1}{1 - b_0} \left(1 + \sum_{i=1}^{i=N} a_i z^{-i} \right) \quad (71)$$

Note that for $C(z)$ to be stable, $H_{control}(z)$ must be minimum phase, which makes developing linear phase FIR filters hard. However, it is theoretically okay for $C(z)$ to be unstable as long as the feedback loop around it is always closed robustly.

$$H_{control}(1) = \frac{1}{1 - b_0} \left(1 + \sum_{i=1}^{i=N} a_i \right) \quad (72)$$

For the controlled waveguide to be stable, we must have the response at DC $H_{control}(1) < 1$. In many cases, we would want $H_{control}(z)$ to be a low pass filter, where often $a_i > 0$ for all i . In order to hold to this requirement, we must have $b_0 \neq 0$. b_0 must be implemented in analog to avoid any delay. Ideally b_0 would

⁷Note that here we assume that we are measuring and actuating exactly the same quantity. In practice, this is not actually possible because all systems are bandwidth limited. In addition, traveling wave-based estimators have additional phase lag associated with them [8]. Nevertheless, we continue with the discussion here because the result is interesting, and we will show that other factors will prevent this type of controller from being easily implemented in the laboratory anyway.




Figure 42: Configuration for controlling the termination's transfer reflection function

be implemented in analog but nevertheless be precisely digitally controllable, i.e. different filter designs $H_{control}(z)$ will generally call for different b_0 .

Besides the preceding complications, obtaining a truly collocated traveling wave measurement/actuation is difficult. Arrays of sensors and actuators must be used [8], and the beamforming algorithms for converting between traveling waves and physically measurable quantities must be implemented without delay, which is difficult. The only possible advantage to this method is that negative feedback can be applied to linearize nonlinear effects. This sort of feedback is typically employed in operational amplifier circuits for this very reason [66].

Since arrays of sensors and actuators are required for general traveling-wave based control, it is desirable to consider formulations involving fewer components, which is possible if knowledge of one of the waveguide terminations is used. We consider this case in the following section.

2.3.3 Termination-Based Control



We now consider termination-based control involving one sensor and one actuator. This control method is attractive in that it allows specification of the termination reflection function. It allows a physical waveguide to be controlled according to linear digital waveguide virtual instrument synthesis techniques [39].

The sensor and actuator do not need to be collocated, although we will see that collocation has desirable properties due to positive realness (see Section 1.10). Consider the control configuration for a one dimensional waveguide as shown in Figure 42, 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 42, 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 $R(f) + A(f)$. $L(f)$ 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 42).

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.

$$L(f) = A(f) + e^{-j4\pi f(\tau_A + \tau_S)}(R(f) + A(f))T(f) \quad (73)$$

$A(f)$ 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 $T(f)$.

$$A(f) = K(f)(R(f) + A(f))(e^{-j2\pi f\tau_A} + e^{-j2\pi f(2\tau_S + \tau_A)}T(f)) \quad (74)$$

Solving (74) for $A(f)$ we arrive at the following:

$$A(f) = K(f)R(f)e^{-j2\pi f\tau_A} \frac{1 + e^{-j4\pi f\tau_S}T(f)}{1 - K(f)e^{-j2\pi f\tau_A}(1 + e^{-j4\pi f\tau_S}T(f))}. \quad (75)$$

Now we can solve for the induced reflection due to the active termination by substituting (75) into (73) and solving for $L(f)/R(f)$.

$$\frac{L(f)}{R(f)} = \frac{K(f)e^{-j2\pi f\tau_A}(1 + e^{-j4\pi f\tau_S}T(f))(1 + e^{-j4\pi f(\tau_A + \tau_S)}T(f))}{1 - K(f)e^{-j2\pi f\tau_A}(1 + e^{-j2\pi f2\tau_S})} + e^{-j4\pi f(\tau_A + \tau_S)}T(f) \quad (76)$$

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 [51]. 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 (76) for $K(f)$, and we discover that $K(f)$ can be noncausal if we do not choose $L(f)/R(f)$ carefully.

$$K(f) = e^{j2\pi f\tau_A} \frac{\frac{L(f)}{R(f)} - e^{-j4\pi f(\tau_A + \tau_S)}T(f)}{(1 + \frac{L(f)}{R(f)})(1 + e^{-j4\pi f\tau_S}T(f))} \quad (77)$$

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:

$$\Lambda(f)e^{-j2\pi f\tau_A} \triangleq \frac{L(f)}{R(f)} \quad (78)$$

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

$$K(f) = \frac{\Lambda(f) - e^{-j2\pi f(\tau_A + 2\tau_S)}T(f)}{(1 + e^{-j2\pi f\tau_A}\Lambda(f))(1 + e^{-j4\pi f\tau_S}T(f))} \quad (79)$$

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 $\Lambda_{damping}(f) = 0$.

$$K_{damping}(f) = \frac{-e^{-j2\pi f(\tau_A + 2\tau_S)}T(f)}{1 + e^{-j4\pi f\tau_S}T(f)} \quad (80)$$

For a quite rigidly-terminated string in the laboratory, $T(f) \approx -0.999$ (see Table 2), but we see that this will cause the poles of the controller $K(f)$ to be dangerously close to the $j\omega$ axis (see (80)). Therefore, **for the controller filter to have well-damped damping poles, we should have something more similar to $T(f) \approx 0.9$** . 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 [5].

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




Figure 43: $|K_{inverting}(f)|$ for $\tau_A = 0$ (collocated, thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)

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. $T(f) = 0$). (79) then implies the following:

$$K_{inverting}(f) = \frac{\Lambda_{inverting}}{1 + e^{-j2\pi f\tau_A} \Lambda_{inverting}}. \quad (81)$$

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 2). It follows then that $K(0) = -99$. Figure 43 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 2.1.1).** Figure 43 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 1.10 to $-K_{inverting}(s/2\pi j)$.

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 $\tau = 0.5$ ms. We still retain $T(f) = 0$.

$$K_{inverting,f0}(f) = \frac{\Lambda_{inverting,f0} e^{-j2\pi f\tau}}{1 + e^{-j2\pi f(\tau_A + \tau)} \Lambda_{inverting,f0}} \quad (82)$$

The result is plotted in Figure 44 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 $|K(0)|$, so the controller is fairly well conditioned, but the phase lag from the term $e^{-j2\pi f\tau}$ prevents the controller from being passive. **Thus, when the waveguide fundamental frequency is adjustable over a large range, the controller is not be passive 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 $T(f) = -0.9e^{-j2\pi f\tau_S}$, where $\tau_S = 1$ 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 45, 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




Figure 44: $|K_{inverting,f_0}(f)|$ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)




Figure 45: $|K_{inv,practical}(f)|$ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)




Figure 46: $|K_{noninvert}(f)|$ for $\tau_A = 0$ (thick red) and $\tau_A \in (0, 0.25]$ ms (thin blue)

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 (79), 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 termination is not positive real.**

SUMMARY

Termination-based control can theoretically be applied for creating an active string termination (see Figure 42). 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 à la PID control.
4. Inducing an inverting termination over a range of positions, and hence making the fundamental frequency adjustable by 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 1.10). 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 (see Section 2.3.4).

2.3.4 Loop gain limitations in the laboratory



From the previous section, we see that some formulations of termination-based




Figure 47: Dashpot attached to a rigidly-terminated semi-infinite vibrating string




Figure 48: Equivalent mechanical configuration to Figure 47 where the impedance of the right-hand string segment is lumped into the dashpot

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 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 47.⁸ Since we are considering velocity waves, the velocity traveling wave components reflect off of the rigid termination with reflection coefficient -1 (see Figure 47). We assume that the dashpot is implemented using feedback control. The dashpot has coefficient large enough to cause a 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 47. The impedances sum because the string endpoint and the dashpot share the same velocity. For velocity waves, the reflection coefficient g is the ratio of the impedance step and the impedance sum [39]:

⁸The 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 controllers. Other analogs exist for other waveguide-based instruments.

$$g = \frac{R_0 - (R_C + R_0)}{R_0 + (R_C + R_0)} = -\frac{R_C}{R_C + 2R_0}. \quad (83)$$

We then solve for R_C .

$$R_C = -\frac{2R_0g}{1+g} \quad (84)$$

By substituting in for g using (53), we arrive at the expression (85) for $R_{C,change\ f_0}$ as a function of the fundamental frequency f_0 of the segment and the decay time constant τ . 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 $\tau = 0.33$ $f_0 = 100\text{Hz}$, $R_{C,change\ f_0} = 65R_0$. Informally speaking, this is longer than a pizzicato note, but shorter than a normal note on a guitar (see Table 2).

$$R_{C,change\ f_0} \triangleq R_C = \frac{2R_0e^{-\frac{1}{f_0\tau}}}{1 - e^{-\frac{1}{f_0\tau}}} \quad (85)$$

DAMPING

To consider pure damping using a dashpot, we need to consider a right termination as well (see Figure 49). g is the reflection coefficient off of the dashpot, so a wave approaching the dashpot continues on to the other segment with gain $(g + 1)$ [41]. 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

$$g_{eff} = -(g + 1)^2. \quad (86)$$

$$g = \sqrt{-g_{eff}} - 1 \quad (87)$$

Combining (87) and (53), we find that

$$R_{C,damped} = -\frac{2R_0g}{1+g} = -\frac{2R_0(\sqrt{-g_{eff}} - 1)}{\sqrt{-g_{eff}}} = -\frac{2R_0(e^{-\frac{1}{2f_0\tau}} - 1)}{e^{-\frac{1}{2f_0\tau}}}. \quad (88)$$



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 $\tau = 0.09$. This implies that $R_{C,damped} \approx 0.11R_0 \ll 65R_0 = R_{C,change\ f_0}$ for $f_0 = 100$.

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 (89), is plotted for $\tau_d = 0.09$ and $\tau_f = 0.33$ in Figure 50. $K_{increase}$ 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. $K_{increase}(6\text{Hz}) \approx 1$ as we could more easily shift the fundamental frequency at low frequencies; however, it is difficult to construct a vibrating string with such low fundamental frequency without having an extraordinarily long string [39]. Moreover, playing such an instrument would not be particularly exciting because the lowest modes would be inaudible [25].




Figure 49: Dashpot attached to a rigidly-terminated vibrating string with reflection coefficients labeled for damping analysis




Figure 50: Increase in loop gain needed for changing the pitch rather than the damping as a function of frequency f_0 ($\tau_d = 0.09$ and $\tau_f = 0.33$)

$$K_{increase}(f_0) = \frac{R_{C,change f_0}(f_0)}{R_{C,damped}(f_0)} = \frac{\frac{2R_0 e^{-\frac{1}{f_0 \tau_f}}}{1 - e^{-\frac{1}{f_0 \tau_f}}}}{\frac{2R_0 (e^{-\frac{1}{2f_0 \tau_d}} - 1)}{e^{-\frac{1}{2f_0 \tau_d}}}} \quad (89)$$



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 $K_{increase} \approx 50\text{dB}$.

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 B.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 also evidences phase lag [52]. To ensure that the control system is stable, the magnitude of the gain around the control loop needs to be less than 1 to ensure stability as soon as the bandwidth limitation causes the phase lag to increase further and further. Consequently, the control loop gain needs to be sufficiently rolled off by the bandwidth limit.⁹ Staying true to our positive real controller philosophy, we are only allowed to use 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 violating one of the characteristics of positive real controllers. 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 $R_C \approx 0.1R_0$.¹⁰

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 $4\text{kHz}/K_{increase} \approx 13\text{Hz}$. Clearly this bandwidth is too small to be musically useful, or even achievable in the laboratory for short strings made of the usual materials.

2.3.5 Active Termination



Given the results from Sections 2.3.4 and 2.3.3, we realize we can only hope to change the fundamental frequency 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 2.3.1, controllers of this form attempting to synthesize long string decay times will suffer from stability problems due to sensor nonlinearity.

We now assume that the string has a perfectly-matched termination as shown in Figure 51. 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

⁹Note that positive real controllers are minimum phase. For minimum phase controllers, there exists a one-to-one relationship between the magnitude response and the phase response [52].

¹⁰We believe that inductive coupling between the actuator and sensor is responsible for this limit (see Section B.4).




Figure 51: Active termination control using a perfectly-terminated string

nothing in the absence of control [41]. 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 controller can 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:

$$\hat{v}_R(t) = \hat{v}(t) + \frac{1}{R_0} f_F(t - \tau). \quad (90)$$

The realization depicted in Figure 51 is difficult to implement in the laboratory for many reasons. To name a few problems, 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. If the string tension is too small, then the wavespeed c will be so low that no active termination can cause the fundamental frequency to be high enough to be musically useful.

We present an alternate possible realization using a piston-type collocated sensor and actuator as depicted in Figure 52. 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 B.5 for information on this type of motor without the added damper.

Next we close an idealized feedback loop around the piston to ensure that it behaves analogously to the system shown in Figure 51 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.

$$K_B(s) = -ms - k/s + (R_0 - R) \quad (91)$$

Regardless of how we implement the perfect termination, we still have $\frac{\partial \hat{v}_R}{\partial v_R} \doteq \frac{\partial \hat{v}}{\partial v}$ due to the estimation of the traveling wave component in (90). That is, because sensor nonlinearity corrupts the velocity estimate, the traveling velocity wave estimate is corrupted to a similar degree. By the argument as in Section 2.3.1, we see that practical implementations attempting to provide long decay times, as characteristic for plucked string instruments, will suffer from stability problems even for piston-type actuators.




Figure 52: Active termination control using a piston-type collocated sensor and actuator


3 Self-Sustaining Control of Acoustic Instruments

3.1 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 [59]. Finally, there is a huge divide between nonlinear control theory and how nonlinear control can be applied realistically in practice [44]. 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 Section 3.2).

3.1.1 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 [26][70]. For small x , the damping is negative, and for large x the damping is positive.

$$m\ddot{x} + a(x^2 - 1)\dot{x} + Kx = 0. \tag{92}$$

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 [59]. 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.


3.1.2 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.¹¹ McIntyre, Schumacher, and Woodhouse present a model applicable to the clarinet, the bowed string, and the flute families [49]. 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 [58].

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 relating the two dual effort and flow wave variables.¹² Consequently, we consider the role of a nonlinear damper when coupled to a mass and a spring in the following section.

3.2 Nonlinear PID Control



Nonlinear proportional-integral-derivative (PID) control may be applied to change the dynamics of an acoustic musical instrument. We first consider a fairly general class of controllers, which result in stable behavior when applied to the resonance model from Section 1.7. We choose the control input u to be the following:

$$u = P_{DD}(\ddot{x})\ddot{x} + P_D(\dot{x}, x)\dot{x} + P_P(\dot{x}, x)x. \quad (93)$$

Then substituting into (1), we obtain a new system with a nonlinear mass, nonlinear damper, and nonlinear spring.

$$(m + P_{DD}(\ddot{x}))\ddot{x} + (R + P_D(\dot{x}, x))\dot{x} + (k + P_P(\dot{x}, x))x = 0 \quad (94)$$

In the interest of presenting a simple stability proof in Section 3.2.3, we restrict ourselves to nonlinear controllers such that $P_{DD}(\ddot{x}) = 0$ and $P_P(\dot{x}, x) = P_P(x)$ for the remainder of the section. However, other nonlinear network theorems from circuit theory may be applied by analogy to prove the uniqueness and stability of equilibria for many other similar systems [74]. We collect the terms together, defining the functions $R(\cdot, \cdot)$ and $K(\cdot)$ to rewrite the nonlinear dynamics equation as follows:

$$m\ddot{x} + R(\dot{x}, x) + K(x) = 0. \quad (95)$$

¹¹Relatively large displacements and velocities can lead to significant distributed nonlinear behavior in some practical situations [35][26]. 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 [26][67].

¹²Indeed, in Section 2.3.3, we observed that the damper was also the simplest element that could be used for changing the fundamental frequency of a waveguide.





Figure 53: Linear damping (left) and saturating damping (right)

3.2.1 Damper




The term $R(\dot{x}, x)$ in (95) describes damping. Although for the stability proof in Section 3.2.3 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 assume initially that $R(\cdot, \cdot)$ is actually independent of x .

Figure 53 (left) shows the characteristic curve describing basic linear damping. For relatively small velocities, viscous damping tends to be linear [29]. Figure 53 (right) shows the characteristic curve describing saturating damping, which could for example describe Coulomb friction in brushed motors [22].

Damping is *locally passive* if $\dot{x}R(\dot{x}, x) \geq 0$ for all $x \in [-x_0 \ x_0]$ and $\dot{x} \in [-\dot{x}_0 \ \dot{x}_0]$ [74]. Damping is globally 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 passive damping may only absorb and never source the power $\dot{x}R(\dot{x}, x)$.

Damping is *strictly locally 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 passive if $x_0 \rightarrow \infty$ and $\dot{x} \rightarrow \infty$. In other words, the damping must be passive and have no deadband. The linear and saturating nonlinear characteristics shown in Figure 53 are both strictly passive because the curves traverse only the first and third quadrants. Figure 54 shows a characteristic curve simulating the effect of holding a stationary bow against a vibrating string [49]. 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 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 [49].

3.2.2 Spring



The term $K(x)$ in (95) describes a nonlinear spring. A linear spring behaves according to $K(x) = kx$ for some constant k . Figure 55 depicts a characteristic curve for a stiffening spring, which is quite common. At low frequencies, loudspeaker drivers exhibit stiffening spring behavior [45]. 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 [35]. 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 [26], p. 149.

For convenience and conforming to the definition in some of the literature [59][74], we define passivity of the spring analogously to the passivity of the damper. $K(x)$ is *locally passive* if $xK(x) \geq 0$ for $x \in [-x_0 \ x_0]$,




Figure 54: Damping characteristic curve due to holding a stationary bow held against a string





Figure 55: Stiffening spring characteristic curve

and it is *strictly locally 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 55 is strictly passive.

3.2.3 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 (95) if the dashpot and spring are strictly locally passive and $|K(x)|$ is bounded locally about $x = 0$. Related proofs for more complicated systems are known from circuit theory [74].

First note that $(x, \dot{x}) = (0, 0)$ is indeed an equilibrium point because both sides of equation (95) become zero for $(x, \dot{x}) = (0, 0)$. Using the analog of a Lyapunov proof for a nonlinear electrical circuit, we define the function $V(x, \dot{x})$ as follows [59][74]:

$$V(x, \dot{x}) = \frac{1}{m} \int_0^x K(\sigma) d\sigma + \frac{1}{2} \dot{x}^2 \quad (96)$$

Because $K(x)$ is strictly locally passive, $\int_0^x K(\sigma) d\sigma$ is locally positive definite about $x = 0$. Furthermore, because $|K(x)|$ is time-invariant and locally bounded and $K(x)$ is locally passive, $K(x)$ is also locally decrescent in x . $\frac{1}{2} \dot{x}^2$ is positive definite and decrescent in \dot{x} . It follows that $V(x, \dot{x})$ is locally positive definite and locally decrescent [59]. Taking the time derivative of (96), we obtain \dot{V} :


$$\dot{V}(x, \dot{x}) = -\frac{R(\dot{x}, x)\dot{x}}{m}. \quad (97)$$

Since the dashpot is strictly locally 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 [59].

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 passive.
- The net spring force is strictly locally passive.
- The initial state is chosen sufficiently close to the origin $(x(t), \dot{x}(t)) = (0, 0)$.

3.2.4 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, 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 3.2.2. The passive nonlinear damper controller causes string vibrations to decay while causing buzzing or rubbing sounds (e.g. see Figure 53, right).





Figure 56: Nonlinear damping control of a vibrating string (top), and string coupled to a nonlinear damper (bottom)

3.3 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 56 reminds the reader of the analogy between controlling a string with a nonlinear damper (Figure 56, top) and a vibrating string coupled to a nonlinear damper at some point along the string (Figure 56, bottom).

To bow the instrument with a virtual bow, we need merely choose the appropriate model of a bow-string nonlinear damping curve. To this end, we move the implicit virtual bow modeled in Figure 54 sideways at velocity v_0 [72] [60]. The result is the horizontally offset characteristic curve shown in Figure 57. 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, the controlled instrument self-oscillates similarly to a bowed string. However, some sets of parameters can lead to period doubling and chaotic behavior [50].

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 58 (left) shows the characteristic curve for a clarinet reed [49]. 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 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 58 (right) shows the instantaneous nonlinear damping characteristic of a trumpet player's lips [58]. 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.

In practice, and especially when building real musical instruments, a sensor and actuator may be effectively




Figure 57: Bowed string nonlinear damping characteristic




Figure 58: Clarinet reed nonlinear damping characteristic (left) and trumpet lip nonlinear damping characteristic (right)




Figure 59: Configuration for event-based control of a vibrating string

collocated for only the lowest modes of vibration (see Section C.4). 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 2.2.4, 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 2.2.4), but $\alpha \ll 2\pi f_0$.

$$F(t) = -R\left(-\alpha \int_0^\infty x(t-\tau)e^{-\alpha\tau}d\tau\right) \quad (98)$$

3.4 Event-Based Control

3.4.1 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 [63], 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 [63]. 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 59.

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 60, 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 that [39]. Figure 60 (a) shows the state of string at rest. Figure 60 (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 60 (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 60 (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 60 (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.¹³

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. 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 [26], and the cycle repeats. The *event* in this scenario is the arrival of the Helmholtz corner.

On a side note, event-based control is very useful in the field of haptics, although in haptics it is traditionally desirable to avoid self-oscillation. Rendering stiff surfaces with feedback control of haptic hardware is difficult and requires expensive hardware. To make hard virtual surfaces feel stiffer, it is often more convenient to apply event-based control than to improve the hardware [47] [46]. In this scenario, stiff surfaces are rendered using the typical penalty-based feedback control method. However, an event-based controller operates in tandem—whenever the haptic device first pushes against a stiff surface, the controller sends a pulse to the haptic device. The pulse not only helps push the device back out of the stiff surface, but it also may contain energy at much higher frequencies than could be easily rendered using the penalty-based feedback control approach. In essence, the event-based controller makes the surfaces *feel* harder.

3.4.2 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.

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

$$\mathcal{F}\{g(t)\} = G(f) = \int_{-\infty}^{\infty} g(t)e^{-2\pi ift} dt. \quad (99)$$

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

$$III(t) \triangleq \sum_{k=-\infty}^{\infty} \delta(t - k) \quad (100)$$

The Fourier transform of the Shah function is the Shah function itself [54].

$$\mathcal{F}\{III(t)\} = III(f) \quad (101)$$

The scaling theorem, as given in (102), is important because it relates time scaling to frequency scaling.

$$\mathcal{F}\{g(ax)\} = \frac{1}{|a|} G\left(\frac{f}{a}\right) \quad (102)$$

¹³In 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 60: Example of traveling velocity waves on a string actuated with a short pulse

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.

$$v(t) = \sum_{k=-\infty}^{\infty} p(t - kT) = p(t) * III\left(\frac{t}{T}\right) \quad (103)$$

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 (102) and the fact that convolution in the time domain corresponds to multiplication in the frequency domain [54].

$$V(f) = P(f) \cdot TIII(fT) \quad (104)$$

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:

$$M(f) = H(f)V(f) = H(f)P(f) \cdot TIII(fT). \quad (105)$$

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 (106). $H(f)P(f)$ is the spectral envelope of $M(f)$. The math from this section is illustrated qualitatively in Figure 61.¹⁴

$$M(f) = H(f)P(f) \cdot \frac{1}{f_0} III\left(\frac{f}{f_0}\right) \quad (106)$$

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 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 [21]. 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 $h(t)$.¹⁵

The spectral envelope of a harmonic signal has a strong influence over the perceived timbre of the signal [25]. This controller is especially useful because it allows great freedom in specifying the spectral envelope of $M(f)$. In contrast, most of the other controllers in this thesis bring about prescribed changes in the spectral envelope, while other controllers bring about complicated changes in the spectral envelope which are hard to predict (see Section 3.5).

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 (106) implies that $M(f) = M_e(f) \cdot \frac{1}{f_0} III\left(\frac{f}{f_0}\right)$. It follows that $M_e(f)$ is indeed the spectral envelope of $M(f)$.

We have implemented this controller in the laboratory for a vibrating string. The interface allows the musician to draw whatever $p(t)$ he or she desires. To make the string "sing", we provide some default pulses $p(t)$ corresponding to different vowels or vocal tract shapes.

¹⁴Figure 61 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 $T \cdot III(fT)$.

¹⁵Note 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 (106), we see that we may achieve the same result of shaping the spectral envelope by specifying $p(t)$ instead.




Figure 61: 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 $P(f)$ is not the exact Fourier transform of $p(t)$.)

3.4.3 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 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 $v_y(t)$ be the velocity measured by the sensor collocated with the actuator. $v_y(t)$ is a combination of the actual string velocity $v(t)$ and measurement noise $n(t)$.

$$v_y(t) = v(t) + n(t) \quad (107)$$

We model $n(t)$ as Gaussian white noise.¹⁶ We then filter the measurement $v_y(t)$ with the filter having impulse response $z(t)$ to obtain $y(t)$.

$$y(t) = z(t) * v_y(t) \quad (108)$$

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 [43]. 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 60). 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 [63]. We have spent some time developing more optimal 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.

¹⁶If $n(t)$ is Gaussian but not white, then we can filter the measurement to make the noise white, and consider instead the goal of detecting the filtered pulse in the filtered sensor signal.

3.5 RMS Level-Tracking Controllers

3.5.1 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 A). 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 [13].

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, etc [39]. 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 straight-forward 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 62 [7]. For example, a simple limiter estimates the RMS level y_p of y using

$$y_p = \sqrt{y^2 * h_{LP}}, \quad (109)$$

where $*$ means convolution and h_{LP} is the impulse response of an (often one-pole) low-pass filter with unity gain at DC.¹⁷ Note that h_{LP} must be non-negative for all time so that the square root in (109) is always defined. The output of the limiter u is calculated as follows for some constant P :

$$u = P \frac{y}{y_p}. \quad (110)$$

It is generally desirable that y_p changes slowly enough that u does not contain too much harmonic distortion [2] (see (140)). However, h_{LP} must allow y_p to change fast enough that the RMS level of u is approximately P . (135) provides a bound on the resonant frequencies of the acoustic musical instrument if the one-pole low-pass filter is applied.

$$u_p = \sqrt{u^2 * h_{LP}} \approx \frac{P}{y_p} \sqrt{y^2 * h_{LP}} = \frac{P}{y_p} y_p = P \quad (111)$$

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

¹⁷In 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.




Figure 62: Control configuration with dynamic range limiter in the feedback loop

RMS level bound r_{RMS} on the input r . 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 $G(s)$.

$$G_{MAX} = \max_{\omega} G(j\omega) \quad (112)$$


G_{MAX} must be bounded because real acoustic musical instruments have some damping at all frequencies [26]. We now form a bound on the RMS level at the output of the instrument x_p :

$$x_p = \sqrt{x^2 * h_{LP}}. \quad (113)$$

It follows that the energy in the musical instrument is bounded assuming that all of the modes are observable at the output. This means that the system is not unstable.

$$x_p \leq G_{MAX}(r_{RMS} + P) < \infty \quad (114)$$

3.5.2 RMS Level-Tracking Controllers



Here we design a more general RMS level-tracking controller that in many 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. 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 [64].

We cannot develop a controller directly for the system in Figure 63 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 64. 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 63, 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

$$u_p = Kx_p \quad (115)$$




Figure 63: RMS level-tracking controller

and

$$u_p = \sqrt{u^2 * h_{LP}} = L\sqrt{y^2 * h_{LP}} = Ly_p \quad (116)$$

by selecting L according to

$$L = \frac{Kx_p}{y_p}. \quad (117)$$

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 (127). L is specified for the adaptive RMS level-tracking controller in (146).

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

$$\dot{x}_p = a_p x_p + b_p \max\{0, u_p\}. \quad (118)$$

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

$$V(e) = e^2/2. \quad (119)$$

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

$$V(\dot{e}) < -\lambda V(e) \quad (120)$$

for some $\lambda > 0$ [59].

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 (116)). 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




Figure 64: RMS level-tracking controller (RMS signals only)

value of x_p , and so the dynamics of the system must decay fast enough alone to satisfy (120). Consider the following where we are taking $e > 0$:

$$\dot{x}_p = a_p x_p \quad (121)$$

$$V(e) = e\dot{e} = e(a_p x_p) < -\lambda e^2/2 \quad (122)$$

$$a_p x_p < -\lambda e/2 = -\lambda(x_p - x_t)/2 \quad (123)$$

$$(a_p + \lambda/2)x_p < \lambda x_t/2. \quad (124)$$

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:

$$V(e) = (x_p - x_t)(\dot{x}_p - \dot{x}_t) = (x_p - x_t)\dot{x}_p = (x_p - x_t)(a_p x_p + b_p u_p). \quad (125)$$

By applying (120) and after some manipulation, we arrive at the following conditions on the controller for exponential convergence with time constant less than $1/\lambda$:

$$\begin{aligned} u_p &> -\frac{\lambda(x_p - x_t)}{2b_p} - \frac{a_p x_p}{b_p} \quad \text{for } x_p < x_t \\ u_p &< -\frac{\lambda(x_p - x_t)}{2b_p} - \frac{a_p x_p}{b_p} \quad \text{for } x_p > x_t \\ u_p &= -\frac{a_p x_p}{b_p} \quad \text{for } x_p = x_t \end{aligned} \quad (126)$$

The light-blue shaded area in Figure 65 (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 65 (right) shows how the valid controller region is restricted as λ is increased from zero and faster convergence is required.

The system trajectories for $u_p = -a_p x_t/b_p$ are shown in Figure 66 (left) [59]. This controller implements the special case of a dynamic range limiter since u_p is held constant [2], so the current work is a generalization of the dynamic range limiter stability result [7]. The light-blue shaded area is drawn to emphasize that some controllers are faster and some are slower (see Figure 66 (left)). The dynamic range limiter can be implemented with equivalent controllers according to (115) and (117):

$$K = \frac{u_p}{x_p} = -\frac{a_p x_t}{b_p x_p} \quad (127)$$




Figure 65: Valid region for convergent controllers (left); Valid region for exponentially-convergent controllers (right)





Figure 66: Trajectories for the dynamic range limiter controller (left); Trajectories for the controller supporting multiple equilibria (right)

$$L = -\frac{a_p x_t}{b_p y_p}. \quad (128)$$

Other more exotic controllers, such as the one whose trajectories are depicted in Figure 66 (right), can support multiple equilibria (in dark blue, Figure 66 (right)).

3.5.4 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 conditions. For the purposes of comparison with the Van der Pol oscillator in section 3.1.1, we assume that the system $G(s)$ describes an undamped mass-spring oscillator. In particular, we have that

$$G(s) = \frac{1}{ms^2 + K}, \quad (129)$$

where the frequency of oscillation $\omega = \sqrt{\frac{K}{m}}$ [26]. 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.

$$H_{LP}(s) = \frac{a}{s + a} \quad (130)$$

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:

$$u = \frac{\zeta}{x_p} x. \quad (131)$$

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.

$$x(t) \approx c_1 \cos(\omega t + \phi_1) \quad (132)$$

$$x^2(t) \approx \frac{c_1^2}{2} (1 + \cos(2\omega t + 2\phi_1)) \quad (133)$$

$$h_{LP}(t) * x^2(t) \approx \frac{c_1^2}{2} \left(1 + \left| \frac{a}{2j\omega + a} \right| \cos(2\omega t + 2\phi_1 + \angle \frac{a}{2j\omega + a}) \right) \quad (134)$$

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.

$$2\omega \gg a \quad (135)$$

$$\frac{1}{x_p} = \frac{1}{\sqrt{h_{LP}(t) * x^2(t)}} \approx \frac{\sqrt{2}}{c_1} \left(1 - \frac{1}{2} \left| \frac{a}{2j\omega + a} \right| \cos(2\omega t + 2\phi_1 + \angle \frac{a}{2j\omega + a}) \right) \quad (136)$$

$$u = \frac{\zeta}{x_p} x \approx \frac{\zeta\sqrt{2}}{c_1} \left(1 - \frac{1}{2} \left| \frac{a}{2j\omega + a} \right| \cos(2\omega t + 2\phi_1 + \angle \frac{a}{2j\omega + a}) \right) (c_1 \cos(\omega t + \phi_1)) \quad (137)$$

$$u \approx \zeta\sqrt{2} \left(\cos(\omega t + \phi_1) - \frac{1}{4} \left| \frac{a}{2j\omega + a} \right| \cos(\omega t + \phi_1 + \angle \frac{a}{2j\omega + a}) - \frac{1}{4} \left| \frac{a}{2j\omega + a} \right| \cos(3\omega t + 3\phi_1 + \angle \frac{a}{2j\omega + a}) \right) \quad (138)$$

Finally, since $2\omega \gg a$, we can simplify the approximation (138):

$$u \approx \zeta\sqrt{2} \left(\cos(\omega t + \phi_1) - \frac{1}{4} \left| \frac{a}{2j\omega + a} \right| \cos(3\omega t + 3\phi_1 + \angle \frac{a}{2j\omega + a}) \right). \quad (139)$$

The approximation, which holds as long as (135) is satisfied, for the harmonic distortion (see Section B.1) introduced by the limiter is


$$k \approx \frac{1}{4} \left| \frac{a}{2j\omega + a} \right|. \quad (140)$$

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 [7]. Some composers also use a clipping nonlinearity instead of a limiter [13].

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 $a/2\pi \approx 16$ Hz [2].

The advantage of the RMS level-tracking controllers over a memoryless controller such as the Van der Pol oscillator explained in Section 3.1.1. 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 [70].

3.6 Adaptive RMS Level-Tracking Controllers



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 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 .

$$V(u_p) = \frac{1}{2}e^2 = \frac{1}{2}(-b_p u_p/a_p - x_t)^2 \quad (141)$$

$$\frac{dV(u_p)}{du_p} = e(-b_p/a_p) \quad (142)$$

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

$$\frac{du_p}{dt} \triangleq -\gamma' \frac{dV(u_p)}{du_p} = -\gamma e \quad (143)$$

The condition $u_p < 0$ needs 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*.

$$u_p = -\gamma \int_{NonNeg} e dt \quad (144)$$

For instance, to implement (144) 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.

$$u_{p,PI} = \max\{0, -\delta e + u_p\} = \max\{0, -\delta e - \gamma \int_{NonNeg} e dt\} \quad (145)$$

Some Wien-Bridge oscillator circuits use this kind of controller to regulate the output amplitude despite any circuit parameter variations [66]. According to (115) and (117), the loop gain L is then

$$L = \frac{u_{p,PI}}{y_p}. \quad (146)$$

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.

$$\frac{X_p(s)}{X_t(s)} = \frac{\delta b_p(s + \gamma/\delta)}{s^2 + (-a_p + b_p\delta)s + \gamma b_p} \quad (147)$$

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

$$\left. \frac{X_p(s)}{X_t(s)} \right|_{s=0} = 1 \quad (148)$$




Figure 67: Ten plucks of linearly increasing magnitude (above), Corresponding loop gain $L(t)$ (below)

3.6.1 Multiple Equilibria Example



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 1.8. Here we use a controller similar to the one depicted in Figure 66 (right) using the velocity feedback effect where

$$y = \dot{x}. \quad (149)$$

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 :

$$L(t) = L_{max} + (L_{min} - L_{max}) \frac{\text{mod}(y_p(t), P_y)}{P_y}. \quad (150)$$

Figure 67 shows the results from the simulation where the string is virtually plucked harder and harder over time. Figure 67 (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. Figure 67 (bottom) reveals that for each pluck, the loop gain L converges to the value 0.08 inducing a stable limit cycle as desired.

3.6.2 Multiple Band-Pass Filter Effect Example



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:

$$y = h_{BP} * x. \quad (151)$$




Figure 68: Driving the RMS of the second harmonic to two values

If multiple band-pass filters are to be applied simultaneously in order to control separate modes, then Figure 63 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 69. Figure 68 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 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.

3.6.3 Resonant Ring Modulation Example



Resonant ring modulation may be implemented using

$$y = x \cdot \cos(2\pi f_c t) \quad (152)$$

for some carrier frequency f_c [7]. 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 $f_c = m f_0$ 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 $2.05 f_0$ are shown in Figure 70. 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 70). 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.




Figure 69: System block diagram for implementing multiple RMS level-tracking controllers concurrently




Figure 70: RMS levels for resonant ring modulation

The companion website provides further explanation of how and why these examples and others were created.¹⁸ 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.

3.7 Summary

We provide a summary in Figure 71 of the characteristics of the nonlinear controllers that we have studied Chapter 3. 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 usually allow the musician to influence the character of the feedback by changing audio effect parameters without affecting the system's stability.

4 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 frequencies of the instrument modes. Note that passive controllers can be applied to all dissipative acoustic instruments, not only to vibrating strings.

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.

A further class of nonlinear controllers works well. 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

¹⁸<http://crma.stanford.edu/~eberdahl/Projects/UnusualDynamics>


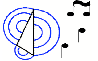

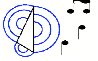
Name	Analysis Class	Linear?	Self-Sustaining?	Passive?	Mechanical Analog	Useful For	Pros	Cons	Implemented?	Simulated?
Passive PID Nonlinear		N	N	Y	Exists	Implementing nonlinear damping and stiffness modulation effects	Stability is guaranteed by passivity. Implements nonlinear modal coupling.	More musically 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 71: Summary of nonlinear controllers covered in Chapter 3




Figure 72: Feedback control of a vibrating string (controller model includes a delay of D seconds)

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 is a huge open area. Musicians may 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.

To help the field progress, we have provided source code for many of our controllers along with an environment for developing and using new controllers (see Appendix A). Many of the controllers have been applied to a professional-quality American Fender Statocaster Deluxe Plus electric guitar. Refer to the website on the feedback guitar for more details:

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

A Toolbox for the Feedback Control of Sound

A.1 Introduction

Feedback control of sound has applications in altering the dynamics of ducts, bars, plates, vibrating strings, and headphones [5]. 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 72. 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 [3].

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 $D \ll \frac{1}{f}$ [27]. 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 [25]. 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 [40].

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 [3].

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 [23]. 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 [3]. 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 [48].

Another alternative we might consider is constructing a low-latency DSP system with a processor supporting native floating point calculations.¹⁹ The most practical configuration for computer music applications involves interfacing a Texas Instruments C6713 DSK with evaluation boards to provide the fast converters. 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.

A.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.²⁰

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

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 [10]. 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()`.

¹⁹These systems are available commercially, but they are prohibitively expensive.

²⁰<http://ccrma.stanford.edu/~eberdahl/Projects/TFCS>




Figure 73: Modular design of TFCS's software components

A.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 [23] (see the dashed lines in Figure 73).

The graphical user interface (GUI) communicates with the control loop thread using open sound control (OSC) [73]. 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) [56]. The default GUI for TFCS is shown in Figure A.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 75). This characteristic makes it possible to implement a wide array of nonlinear dashpots,




Figure 74: The GUI for TFCS provides buttons, toggles, sliders, etc.

nonlinear springs, etc.

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 [27]. Figure A.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.

$$D = \frac{1}{2f_s} + D_{processing} \quad (153)$$

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 A.3, and the delay can be estimated by visual inspection. For our system, we estimate that $D \approx 24\mu s$ for $f_s = 40\text{kHz}$. By applying (153), we find that $D_{processing} \approx 11.5\mu s$.²¹

²¹For our National Instruments PCI6221 data acquisition card, acquisition takes about $7\mu s$, while the DAC has a full-scale




Figure 75: The GUI also allows arbitrary nonlinear characteristic curves for controller elements to be drawn with the mouse.




Figure 76: A continuous time signal (in blue) as approximated by a zero-order hold (dash-dotted in red)




Figure 77: Estimated jitter for TFCS given no additional load ($f_s = 40\text{kHz}$)

A.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 79).

We present some estimations of the jitter observed when using TFCS under various conditions given the hardware listed in Section A.2. TFCS provides real time feedback about the jitter performance.²² In Figure 77, we present the performance of TFCS under fairly optimal conditions: $f_s = 40\text{kHz}$ and no system loading by other applications. Figure 77 (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 77 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 77, bottom).

Even though the control code is isolated on CPU 0, if CPU 1 is heavily loaded, the jitter performance is affected. Figure 78 shows the statistics gathered over several minutes. The conditions were the same as for Figure 77 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 79 correspond to a test under the same conditions as in Figure 77 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.

settling time of roughly $4\mu\text{s}$. The approximate remaining delay may be due to various sources such as acquiring the PCI bus [23].

²²This feature could be disabled to make it possible to do more computations during the control loop.




Figure 78: Estimated jitter for TFCS given significant system load ($f_S = 40k\text{Hz}$)




Figure 79: Estimated jitter for TFCS given a system loaded by constantly dragging a window back and forth ($f_S = 40k\text{Hz}$, 0.8% overruns)




Figure 80: Estimated jitter for TFCS with no load ($f_S = 50kHz$, 4.6% overruns)

Finally we tested whether we could increase the sampling rate further. We found $f_S = 50kHz$ to roughly be the maximum rate (see Figure 80). Here we were not running the full TFCS control loop, rather the control loop merely read an input 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.

A.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

A.6 Initialization Files

A.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/LogVol100 isolcpus=1 noirqbalance
  initrd /initrd-2.6.19.img
```

A.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
```

A.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
```

A.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
```




Figure 81: SISO system

A.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)
```

B 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.

B.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 81. Let the input $u = A \cos \omega t$. Most deterministic systems²³ will produce output of the form

$$y = \sum_{n=0, \dots, \infty} B_n \cos(n\omega t + \phi_n). \quad (154)$$

We define the total harmonic distortion k to be

$$k = \sqrt{\frac{\sum_{n=2, \dots, \infty} \frac{1}{2} B_n^2}{\sum_{n=1, \dots, \infty} \frac{1}{2} B_n^2}} \quad (155)$$

and the partial harmonic distortion components k_n to be

$$k_n = \left| \frac{B_n}{B_1} \right| \quad (156)$$

for $n = 2, \dots, \infty$ [66].

²³Exceptions include chaotic systems and systems exhibiting subharmonic behavior.




Figure 82: Lorentz' law type actuator

B.2 Actuators

In actuator design, there is typically a trade-off between nonlinearity and maximum displacement [45]. 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.

B.2.1 Lorentz' Law Type

The Lorentz' law type of string actuator behaves simply. Figure 82 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 82). 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.²⁴

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 (157). According to the geometry shown in Figure 82, 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 [45] [27].

$$\mathbf{F}_{\text{string}} = L_A \mathbf{I} \times \mathbf{B} \quad (157)$$

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 loop.²⁵

²⁴Some 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.

²⁵Of 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 83: Lorentz' law type actuator (viewed from above)




Figure 84: Lorentz' law type actuator (oblique angle)

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 regular piece of 22 AWG solid wire, which was 0.644mm in diameter. The insulation was removed. 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 83 and 84.

B.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 [36] and Sustainiac [37] electric guitar string vibration sustainers. A hand-wound reluctance force actuator is shown in Figure 85. It consists of a coil wound around a magnetically permeable core, for instance iron. In Figure 85, the core is a pole piece removed from a Fender electric guitar pickup. 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 [61]. 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.²⁶ In general, f_X and $i(t)$ will be nonlinearly related, as

²⁶The remainder of Section B.2.2 is taken almost verbatim from a paper written by Edgar Berdahl, Steven Backer, and Julius




Figure 85: Hand-wound reluctance force actuator

depicted in Figure 86 (right). f_X will be approximately proportional to $i^2(t)$, except for fields so large that the string saturates magnetically, in which case f_X becomes nearly linearly related to i (circle in Figure 86). 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 86, left). This is undesirable. We have made measurements showing that when string displacements are relatively small, the system is roughly time-invariant [6]. Unfortunately our analysis is not so simple as in [61] 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 86: 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 $\mathbf{B}_S(t)$.

$$\mathbf{B}(t) = \mathbf{B}_L(t) + \mathbf{B}_S(t) \quad (158)$$

$\mathbf{B}_L(t)$ 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 $i(t)$ [41]. $\mathbf{B}_L(t)$ 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.

$$\mathbf{B}_L(t) = \mathbf{B}_{pm} + \mathbf{B}_{coil}(t) \approx \mathbf{B}_{pm} \quad (159)$$

The qualitative plot in Figure 87 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

O. Smith III [6].

wider core and multiple permanent magnets as compared to the geometry in Figure 85, 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 [41].




Figure 87: 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.²⁷ 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 $i(t)$ 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.

$$\Delta\mathbf{B}(t) \triangleq \mathbf{B}_S(t) = \mathbf{B}_{\text{coil}}(t) \quad (160)$$

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.

$$U(t) = -\mathbf{m} \cdot \mathbf{B}_{\text{coil}}(t) = -\cos(\theta)|\mathbf{m}||\mathbf{B}_{\text{coil}}(t)| \quad (161)$$

The force in the x -direction can be determined by taking the derivative with respect to the x -direction.

$$f_X = -\frac{\partial U(t)}{\partial x} \quad (162)$$

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 86:

$$f_X \propto x^{-M-1} \quad (163)$$

and

$$f_X \propto i(t). \quad (164)$$

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

²⁷For 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.



Figure 88: Photodiode type optical sensor in the laboratory (LED not shown)

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 [36].

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 (165), 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.

$$\frac{V(s)}{I(s)} = \frac{1}{L} \cdot \frac{1}{s + \frac{R}{L}} \quad (165)$$

B.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.

B.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 88, which is recommended by Weinreich and Caussé [72]. 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 89, 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 related to the motion of the string in the plane parallel to the sensitive surface.



Figure 89: Photodiode type optical displacement sensor

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 88). 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.²⁸ 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 [1]. 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 $f_{3dB,theoretical} = 21.5kHz$.²⁹

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 [15]. 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 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 [15]. For this particular LED, the radiant intensity falls off by only 6% by 10 degrees from the normal [68]. 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.

²⁸Christine 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.”

²⁹We 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 [1], but we considered $f_{3dB,measured} \approx 8.5kHz$ to be large enough for our application.



Figure 90: Halogen lamp



Figure 91: Nonuniform radiation pattern of a coiled halogen lamp

Halogen lamps are worth considering because they produce so much output power. For instance, consider the 20W halogen lamp shown in Figure 90. The lamp is made by JC Halogen Bulb China, is rated at 12V, and has a BI-PIN base. 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 91 by holding a white sheet of paper behind the lamp.

Nevertheless, Figure 92 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 92. We found that using this light source in conjunction with either of the photodiodes previously discussed greatly improved the 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 2 and 3, 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 [34]. Because the phototransistor produces such a large voltage output, the measurement is likely less susceptible to inductive coupling effects (see Section B.4).




Figure 92: More uniform radiation pattern of lamp

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 [34]. We damped the string by wrapping it in felt to reduce 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 93 [11]. 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 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.

B.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 [4]. 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.³⁰ Given that such piezoelectric elements typically have small mass and large stiffness, their mechanical resonance frequency is above the hearing threshold [4]. Hence, the force exerted on the piezoelectric element is proportional to the string’s

³⁰For 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 93: Piezoelectric guitar string termination sensor from Graphtech



Figure 94: Piezoelectric guitar string termination sensor from Graphtech

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 [4]. A piezoelectric guitar string termination sensor from Graphtech for a Fender electric guitar is shown in Figure 94. We have used this type of sensor for the actively controllable electric guitar detailed in Section C.4.

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 [28].




Figure 95: Spectrum analyzer showing the nonlinearity in the sensor signal for a S2387-1010R silicon photodiode illuminated by the halogen lamp

B.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 B.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 [35]. 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.³¹ 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 95, 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:

³¹For the preferred configuration with the PD638B photodiode and IR204/H16/L10 LED in Section B.3.1, which was used for experiments in Chapters 2 and 3, we estimate that $k_{displacement} \approx -40\text{dB}$.




Figure 96: Collocation geometry

$$k_{displacement} \approx 10 \log_{10}(10^{-35/10} + 10^{-39/10} + 10^{-42/10} + 10^{-49/10}) \approx -33\text{dB}. \quad (166)$$

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 (167).

$$k_{velocity} \approx 10 \log_{10}(10^{(-35+6)/10} + 10^{(-39+6 \log_2 \frac{3}{2})/10} + 10^{(-42+12)/10} + 10^{(-49+6 \log_2 \frac{5}{2})/10}) \quad (167)$$

We see that the THD of the velocity signal is even larger than the THD of the displacement signal.

$$k_{velocity} \approx -26\text{dB} \quad (168)$$

B.4 Collocation Considerations

As explained in Section 1.10, we need to be able to collocate sensors and actuators. We have chosen to collocate a photodiode optical displacement sensor (see Section B.3.1) with a Lorentz' Law type actuator (see Section B.2.1). We collocate the elements in the presence of the string as shown in Figure 96. 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 96). Section 2.1.9 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 88), 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 implies that the current I flowing along the string induces a voltage V_{as} at the input to the operational amplifier [41].

$$V_{as} = L_{as}(x) \frac{\partial I}{\partial t} \quad (169)$$

Because of the differentiation in (169), 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.




Figure 97: String termination actuator concept

B.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 we did not pursue the design further. The general concept behind the active string termination is shown in Figure 97. 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.

C Musical Instruments

C.1 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.³² 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” [36] 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 in the same sense; instead, we transmit sounds into the piano, which then emanate acoustically

³²This section is taken almost verbatim from a paper written by Edgar Berdahl, Steven Backer, and Julius O. Smith III [6].



Figure 98: Close up of the motors

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 Caus e [72] 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 98). 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 B.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,³³ though they are by no means representative of the full capabilities of the

³³<http://crma.stanford.edu/~sbacker/empp/>

instrument.

In the block diagram shown below (see Figure 99), the basic structure of the device is depicted for one channel.³⁴ 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 100). The electromagnets themselves were originally intended for holding applications.




Figure 99: System block diagram for one channel

³⁴A total of 12 channels / notes were implemented.



Figure 100: Electromagnetically prepared piano prototype

C.2 Monochord




Figure 101: Monochord (tuning peg end)




Figure 102: monochord (control end)

C.3 One String Guitar



Figure 103: One String Guitar




Figure 104: One string guitar (close up)

C.4 Feedback Guitar



Figure 105: Feedback guitar

References

- [1] *Photodiode Amplifiers: OP AMP Solutions*. McGraw-Hill Professional, Columbus, OH, 1995.
- [2] 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*, number Preprint 5914, 2003.
- [3] 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.
- [4] Antonio Arnau. *Piezoelectric transducers and applications*. Springer, New York, NY, 2004.
- [5] Volker Bartels. Headset with active noise-reduction system for mobile applications. *Journal of the Audio Engineering Society*, 40(4):277–281, 1992.
- [6] 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.
- [7] Edgar Berdahl and Julius O. Smith III. Some physical audio effects. In *Proceedings of the 9th International Conference on Digital Audio Effects*, pages pp. 165–168, Sept. 18-20 2006.
- [8] Edgar Berdahl and Julius O. Smith III. Estimating the state of a one-dimensional waveguide. In *153rd Meeting of the Acoustical Society of America*, June 4-7 2007.
- [9] Edgar Berdahl and Julius O. Smith III. Pid control of a vibrating string. In *153rd Meeting of the Acoustical Society of America*, June 4-7 2007.

- [10] Edgar Berdahl, Nelson Lee, Günter Niemeyer, and Julius O. Smith III. Estimating the state of a one-dimensional waveguide. In *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.
- [11] Edgar Berdahl and Julius Smith. Transfer function measurement toolbox.
- [12] 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.
- [13] 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.
- [14] 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*, 1999.
- [15] 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.
- [16] Henri Boutin. Contr le actif sur instruments acoustiques. ATIAM Master's Thesis, Sept. 2007.
- [17] 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.
- [18] Henri Boutin and Charles Besnainou. Physical parameters of the violin bridge changed by active control. In *Proc. of the Acoustics '08 Conference*, June, 2008.
- [19] Claude Cadoz, A. Luciana, and Jean-Loup Florens. Synth se musicale par simulation des m canismes instrumentaux. *Revue d'acoustique*, 59:279–292, 1981.
- [20] Remy Chollet, Guillaume Aeberli, and Charles Besnainou. Modifier la resonance de helmholtz de la guitare par contr le actif. In *Proceedings of the 5th French Congress on Acoustics*, pages 248–250, Sept. 3-6 2000.
- [21] Perry Cook. *Identification of Control Parameters in an Articulatory Vocal Tract Model, With Applications to the Synthesis of Singing*. PhD thesis, Stanford University, 1991.
- [22] Nicola Diolaiti, Guenter Niemeyer, Federico Barbagli, Kenneth Salisbury, and Claudio Melchiorri. The effect of quantization and coulomb friction on the stability of haptic rendering. *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 237–246, 2005.
- [23] 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.
- [24] Wang Enping and Wang Chaozhu. Some properties of the positive real matrix. *Acta Mathematicae Applicatae Sinica*, 1(1), 1984.
- [25] Hugo Fastl and Eberhard Zwicker. *Psychoacoustics: Facts and Models*. Springer, New York, NY, 2006.
- [26] Neville Fletcher and Thomas Rossing. *The Physics of Musical Instruments*. Springer, New York, NY, 2nd edition, 1998.
- [27] Gene Franklin, J. Powell, and Abbas Emami-Naeini. *Feedback Control of Dynamic Systems*. Prentice Hall, fifth ed. edition, 2005.
- [28] 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.

- [29] Michael Friswell and J. E. Mottershead. *Finite Element Model Updating in Structural Dynamics*. Kluwer Academic Publishers, Boston, MA, 1995.
- [30] 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.
- [31] Jean Guérard and Xavier Boutillon. Real time acoustic wave separation in a tube. In *in Proceedings of the International Symposium on Musical Acoustics*, number 451, August 19-22 1997.
- [32] Sathya Hanagud and Steven Griffin. Active structural control for a smart guitar. In *Proceedings of the Fourth European Conference On Smart Structures and Materials*, July 6-8 1998.
- [33] Blake Hannaford. A design framework for teleoperators with kinesthetic feedback. *IEEE Transactions on Robotics and Automation*, 5(4):426–434, August 1989.
- [34] Roger Hanson. Optoelectronic detection of string vibration. In *The Physics Teacher*, volume 25, pages 165–166, 1987.
- [35] Roger Hanson, James Anderson, and Kent Macomber. Measurements of nonlinear effects in a driven vibrating wire. *Journal of the Acoustical Society of America*, 96(3):1549–1556, September 1994.
- [36] Gregory Heet. String instrument vibration initiator and sustainer. U.S. Pat. 4,075,921, 1978.
- [37] Alan Hoover. Controls for musical instrument sustainers. U.S. Pat. 6,034,316, 2000.
- [38] Julius O. Smith III. *Introduction to Digital Filters with Audio Applications*. W3K Publishing, <http://ccrma.stanford.edu/~jos/pasp/>, 2007.
- [39] Julius O. Smith III. *Physical Audio Signal Processing: For Virtual Musical Instruments and Audio Effects*. W3K Publishing, <http://ccrma.stanford.edu/~jos/pasp/>, 2007.
- [40] Julius O. Smith III. *Mathematics of the Discrete Fourier Transform (DFT) with Audio Applications*. W3K Publishing, <http://ccrma.stanford.edu/~jos/mdft/>, 2008.
- [41] Umran Inan and Aziz Inan. *Engineering electromagnetics*. Menlo Park, CA, 1999. Addison Wesley Longman.
- [42] J. Douglas Rollow IV. *Active Control of Spectral Detail Radiated by an Air-Loaded Impacted Membrane*. Pennsylvania State University, 2003.
- [43] Nikil Jayant and Peter Noll. *Digital Coding of Waveforms: Principles and Applications to Speech and Video*. Prentice Hall, March 1984.
- [44] Hassan Khalil. *Nonlinear Systems, 3rd Ed.* Prentice Hall, Upper Saddle River, NJ, 2001.
- [45] Wolfgang Klippel. Loudspeaker nonlinearities - causes, parameters, symptoms. In *Proceedings of the 119th Convention of the Audio Engineering Society/Audio Engineering Society*, number Preprint 6584, New York, NY, October, 7-10 2005.
- [46] Dimitrios Kontarinis and Robert Howe. Tactile display of vibratory information in teleoperation and virtual environments. *Presence: Teleoperators and Virtual Environments*, 4(4):387–402, 1995.
- [47] Katherine Kuchenbecker, Jonathan Fiene, and Günter Niemeyer. Improving contact realism through event-based haptic feedback. *IEEE Transactions on Visualization and Computer Graphics*, 12(2):219–230, March/April 2006.
- [48] 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.
- [49] 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.

- [50] Georg Mueller and Werner Lauterborn. The bowed string as a nonlinear dynamical system. *Acustica*, 82:657–664, 1996.
- [51] Philip Nelson and Steve Elliott. *Active Control of Sound*. Academic Press, Elsevier Publishing, Burlington, MA, 1993.
- [52] Alan Oppenheim and Ronald Schaffer. Discrete-time signal processing, 2nd edition. Upper Saddle River, NJ, 1999. Prentice Hall.
- [53] Gary Osborne and Alan Hoover. Sustainer for a musical instrument. U.S. Pat. 5,932,827, 1999.
- [54] Brad Osgood. *Lecture Notes for EE261: The Fourier Transform and its Applications*. Stanford University, 2007.
- [55] Allen Pierce. *Acoustics: An Introduction to Its Physical Principles and Applications*. Acoustical Society of America, June 1989.
- [56] Miller Puckette. Pure data: another integrated computer music environment. In *Proceedings of the Second Intercollege Computer Music Concerts*, pages 37–41, Tachikawa, 1997.
- [57] 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.
- [58] 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.
- [59] Shankar Sastry. *Nonlinear Systems: Analysis, Stability, and Control*. Springer, New York, NY, 2004.
- [60] Stefania Serafin. *The Sound of Friction: Real-Time Models, Playability and Musical Applications*. PhD thesis, Stanford University, Stanford, CA, USA, June 2004.
- [61] 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.
- [62] Tim Stilson. Forward-going wave extraction in acoustic tubes. In *Proceedings of the International Computer Music Conference*, 1995.
- [63] Karl Åström. Event based control. *Analysis and Design of Nonlinear Control Systems*, pages 127–147, 2008.
- [64] Karl Åström and Bjorn Wittenmark. *Adaptive Control*. Addison-Wesley, Boston, MA, second edition, 1995.
- [65] 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.
- [66] Ulrich Tietze and Christoph Schenk. *Halbleiter-Schaltungstechnik*. Springer, Berlin, Germany, 11th edition, 1999.
- [67] 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.
- [68] Jaine Tsai. Technical data sheet 3mm infrared led t-1 ir204/h16/110. Technical report, Everlight Corp., 2005.
- [69] Mac Van Valkenburg. *Introduction to Modern Network Synthesis*. John Wiley and Sons Inc., Hoboken, NJ, 1960.
- [70] Balthasar van der Pol. The nonlinear theory of electric oscillations. In *Proceedings of the Institute of Radio Engineers*, volume 22, pages 1051–1086, 1934.

- [71] 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.
- [72] 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.
- [73] 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.
- [74] Chai Wah Wu. Qualitative analysis of dynamic circuits. *Wiley Encyclopedia of Electrical and Electronics Engineering*, 1999.