VIRTUAL STRINGED INSTRUMENTS

Nelson Lee — Stanford University, CCRMA

Julius O. Smith III — Stanford University, CCRMA

ABSTRACT

This reader contains an overview of physical-modeling methods for stringed instruments. It should be entirely self-contained, including external links to labs and more detailed content on relevant subjects. By thoroughly understanding this reader, you should be well equipped to tackle the ensuing advanced labs in the RealSimple tree.

Contents

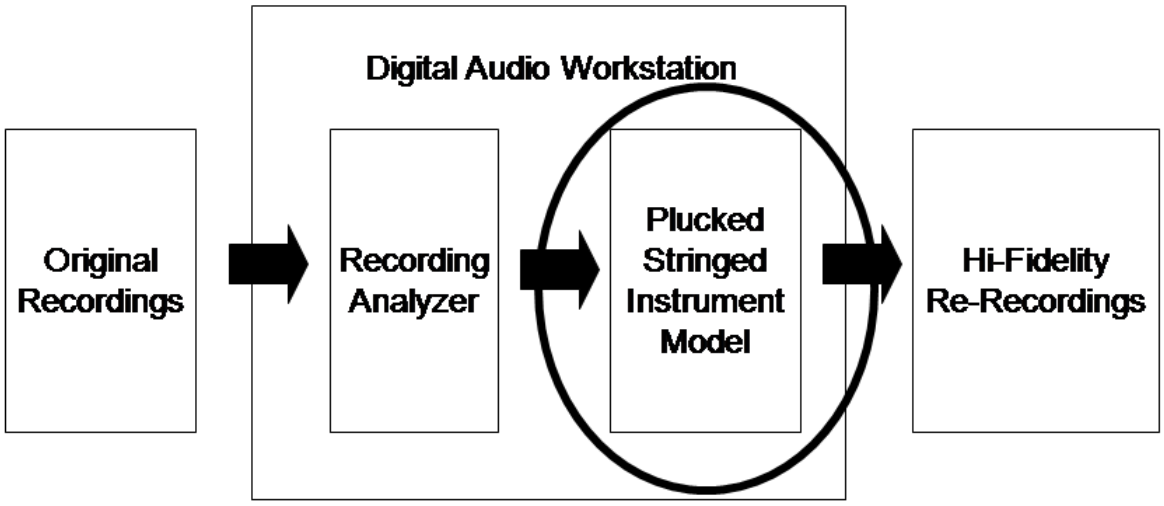

Figure 1. Block diagram of a performance parameter estimation and resynthesis of old recordings application.

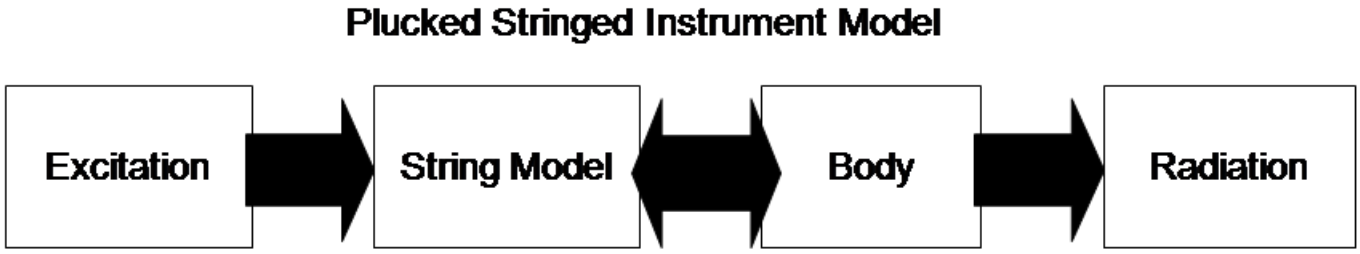

Figure 2. Stringed instrument decomposition abstraction block diagram.

1. INTRODUCTION

In this reader, we discuss methods for real-time synthesis of stringed instruments. Interest in this topic is wide and varying, as both studio and performance uses for realistic virtual stringed instruments are becoming increasingly possible with gains in computing power.

Having a high-fidelity physics-based virtual stringed instrument model is useful for many applications. Current sample-based synthesizers, which are based on audio recordings, do not allow fine control of the excitation of the strings of the instrument. Synthesizers based on physical models, in contrast, promise unlimited control over the expressive nuances of string excitation. The synthesis parameters are also more intuitive, since they have corresponding meanings in the physical world. Changing intuitive parameters can result in more realistic changes to the sound produced.

Another of many applications includes automatic transcription and resynthesis of old recordings. Given that there is a mechanism for processing old recordings and mapping them to how they were played on the instrument, with a high-fidelity synthesis model and the necessary performance parameters, resynthesis of old recordings can be made to sound like what they would, had they been made with today's technology. Figure 1 shows the block diagram of such a system. The focus of this reader is on the third circled block, the physical model of the plucked stringed instrument.

The goal of this reader is to outline procedures for making a virtual stringed musical instrument based on a combination of physical theory and laboratory measurements from a real instrument. Since this topic is too large to be covered in the available space, we will make extensive use of pointers to supporting information. In addition to the traditional bibliographic citations, we will refer the reader to additional online books, related websites, and laboratory exercises covering elementary models and techniques used. Our goal is to make it possible to follow links in this reader in order to flesh out the complete details of the theory and practical techniques summarized here. For the advanced reader (such as a seasoned graduate student working in the area of virtual musical instrument design), this reader will hopefully prove sufficiently self-contained to be used as a laboratory guide.

In the following sections, we briefly review elementary components of stringed musical instruments and how they may be modeled efficiently for real-time digital synthesis applications. We then summarize practical measurement and calibration techniques for various instrument modeling components. Finally, we discuss methods specifically for estimating parameters of a *plucked* stringed instrument.

A useful abstraction that illustrates our modeling approach for virtual stringed instruments is shown in Figure 2. Not only is this decomposition useful for compartmentalizing from a modeling perspective, it is useful in performing measurements on the instrument as well as following the physical flow of how a stringed instrument is played: energy injected into the system to how energy reaches our ears by pressure waves created by the vibration of the instrument's body.

2. ELEMENTS OF STRINGED INSTRUMENTS

In all stringed instruments, the strings are put into motion by an external force typically applied by a finger, plectrum ("pick"), hammer, or bow. The vibrating string transfers energy into the body/resonator of the instrument, which in turn produces pressure waves in the air that propagate to our ears.

2.1. Vibrating String

In this section, we briefly describe the physics of the string and its corresponding virtual representation. We describe how this model of the string is analogous to its physical counterpart and describe various initial-conditions for the different ways the string can be set into motion. We then discuss more complex modeling techniques for the vibrating string,

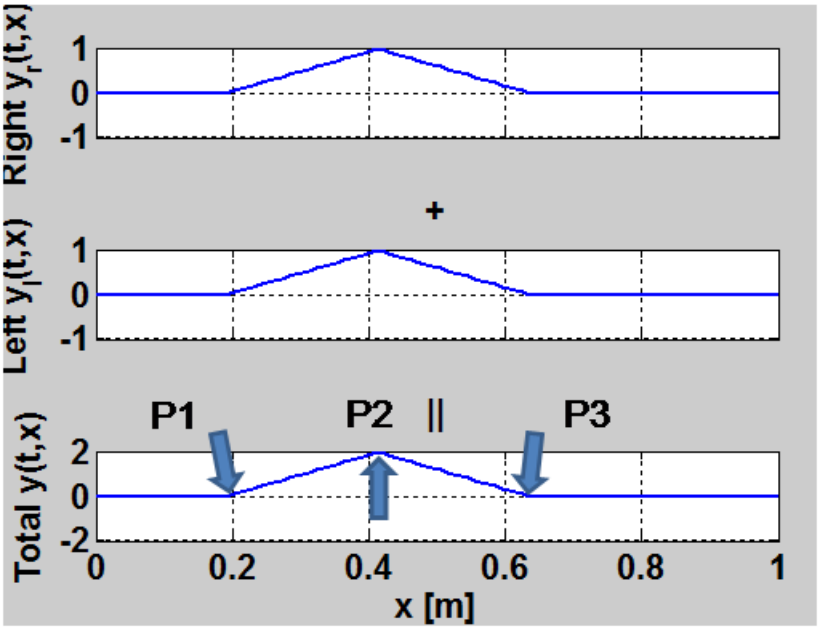

Figure 3. t=0, Figures 3 to 10 show an example of the traveling wave solution to a string that is displaced to an amplitude of 2 at time t=0. We note that the string is held in place at three points, P1, P2 and P3 at t=0. The top plot of each figure shows the right traveling wave, the middle plot the left traveling wave and the bottom plot the physical displacement, equal to the sum of the top and middle plots. *Plots generated using scripts developed by Ed Berdahl.*

initial pitch shifts, the coupling of multiple strings and the vibration of the string in two orthogonal planes and how such modeling techniques can describe natural phenomena such as two-stage decay.

2.1.1. D'Alembert's Wave Equation

The formulation of the Wave Equation and its solution by d'Alembert in 1747 [1] is the theoretical starting-point for physical stringed models. The wave equation, written as

$$K\frac{\partial^2 y}{\partial x^2} = \varepsilon\frac{\partial^2 y}{\partial t^2} \tag{1}$$

where K is the string tension, ε is the linear mass density and $y(t,x)$ is the string displacement as a function of time (t) and position along the string (x). It can be derived directly from Newton's second law applied to a differential string element. In addition to introducing the 1D wave equation, d'Alembert introduced its solution in terms of traveling waves:

$$y(t,x) = y_r(t-\frac{x}{c}) + y_l(t+\frac{x}{c}) \tag{2}$$

where $c = \sqrt{K/\varepsilon}$ denotes the wave propagation speed. Though each individual traveling wave is unobservable in the physical world, we use their constructs for modeling the physical behavior of the string, whose displacement is equal to the sum of the two traveling waves. Figures 3 to 10 show how a physical string's displacement, initially displaced to a triangular pulse, is equal to the sum of its left and right traveling waves at time $t = 0$. The blue arrows in Figure 3 show the points of displacement at time $t = 0$. Subsequent frames are shown, as the bottom plot of each frame corresponds to the physical displacement of the string. The top and bottom plots correspond to the right and left traveling waves, respectively. For a detailed derivation of the solution to the Wave Equation, we refer readers to http://ccrma.stanford.edu/ jos/pasp/Traveling_Wave_Solution_I.html.

D'Alembert's solution to the wave equation serves as the theoretical foundation of which physical models for stringed instruments are based upon.

2.1.2. The Delay Line

In this section we present the fundamental computing block used in physical modeling. We present different views of this data structure and use it for building string models for the remainder of this reader.

The data structure of the delay line naturally arose from digital computing. From a Computer Science perspective, the delay line is simply a vector of values. In the musical acoustics community, the delay line is the digital implementation

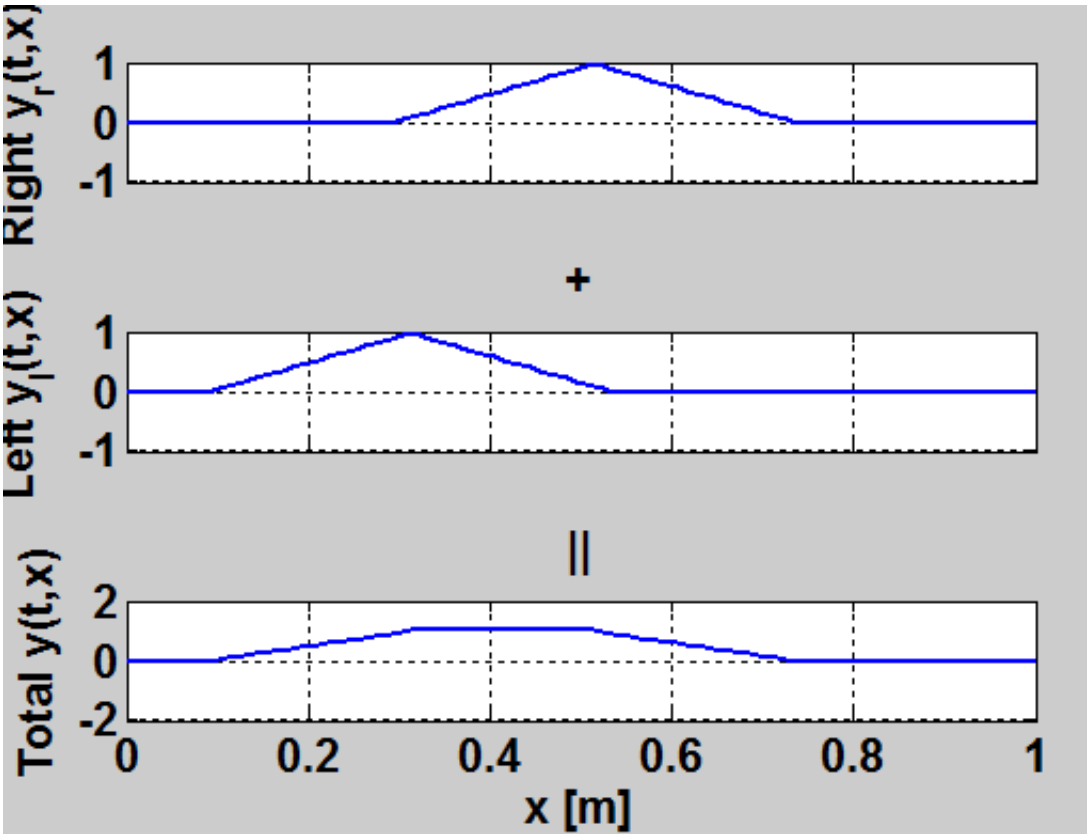

Figure 4. t=5

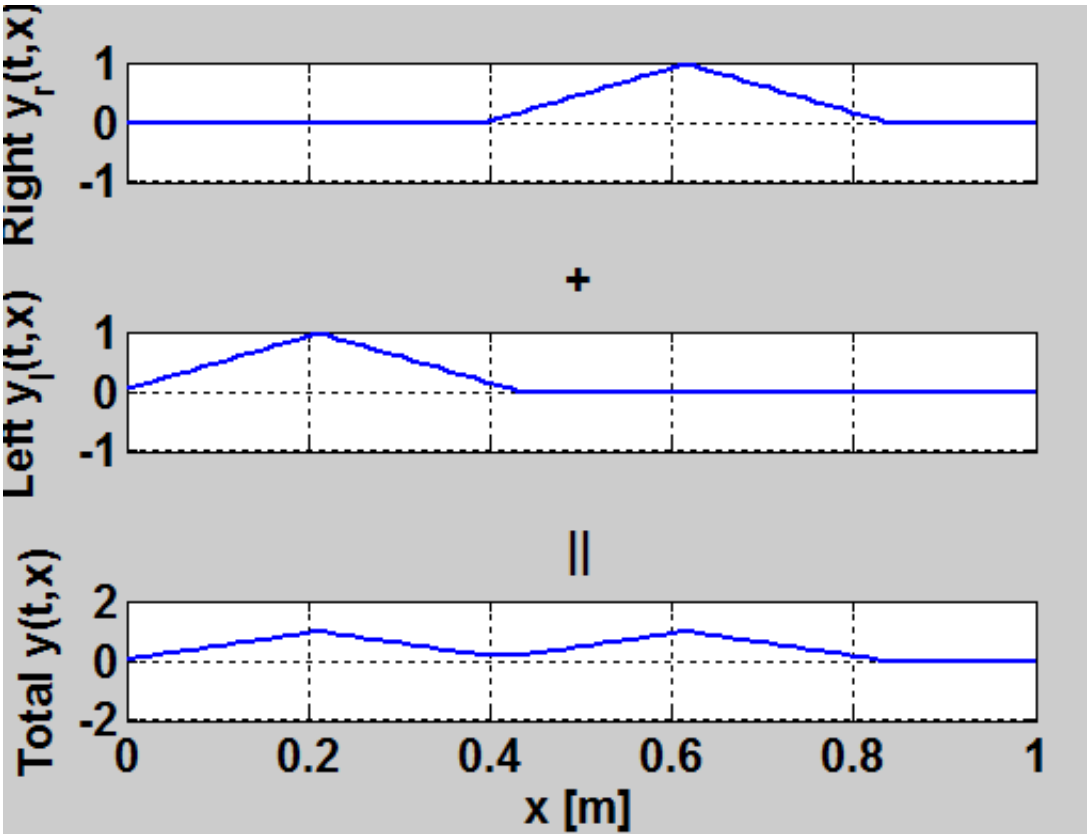

Figure 5. t=10

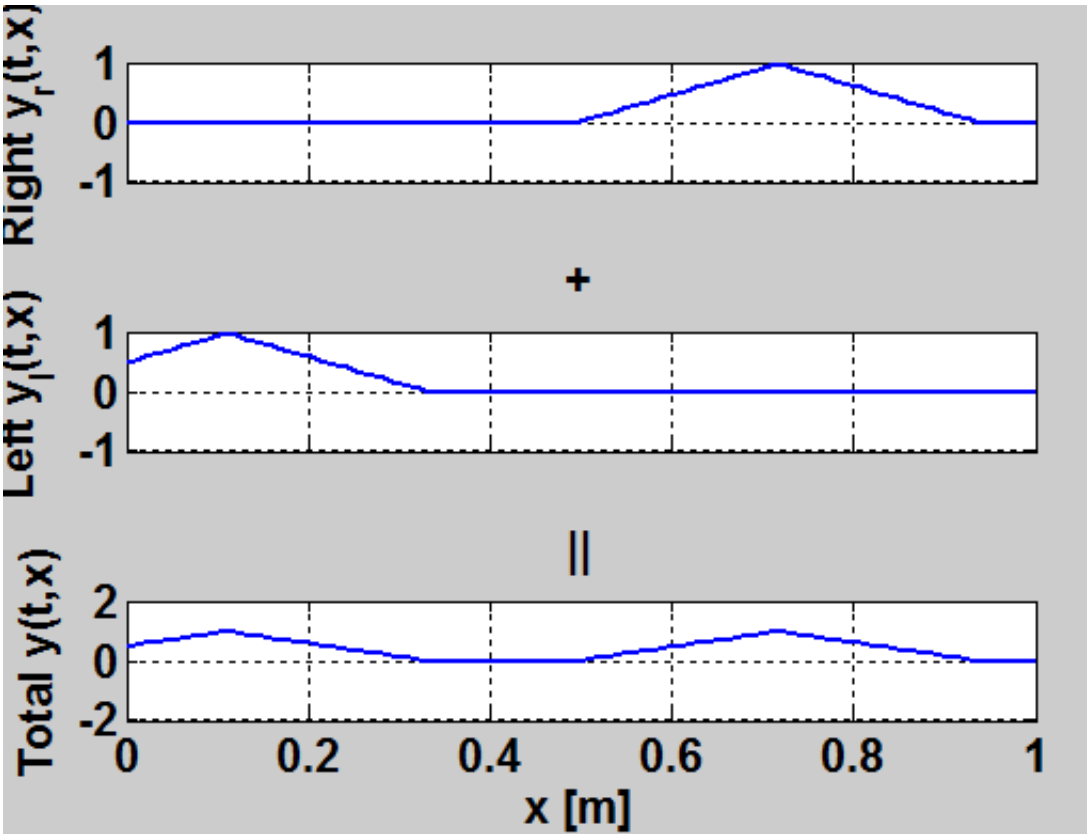

Figure 6. t=15

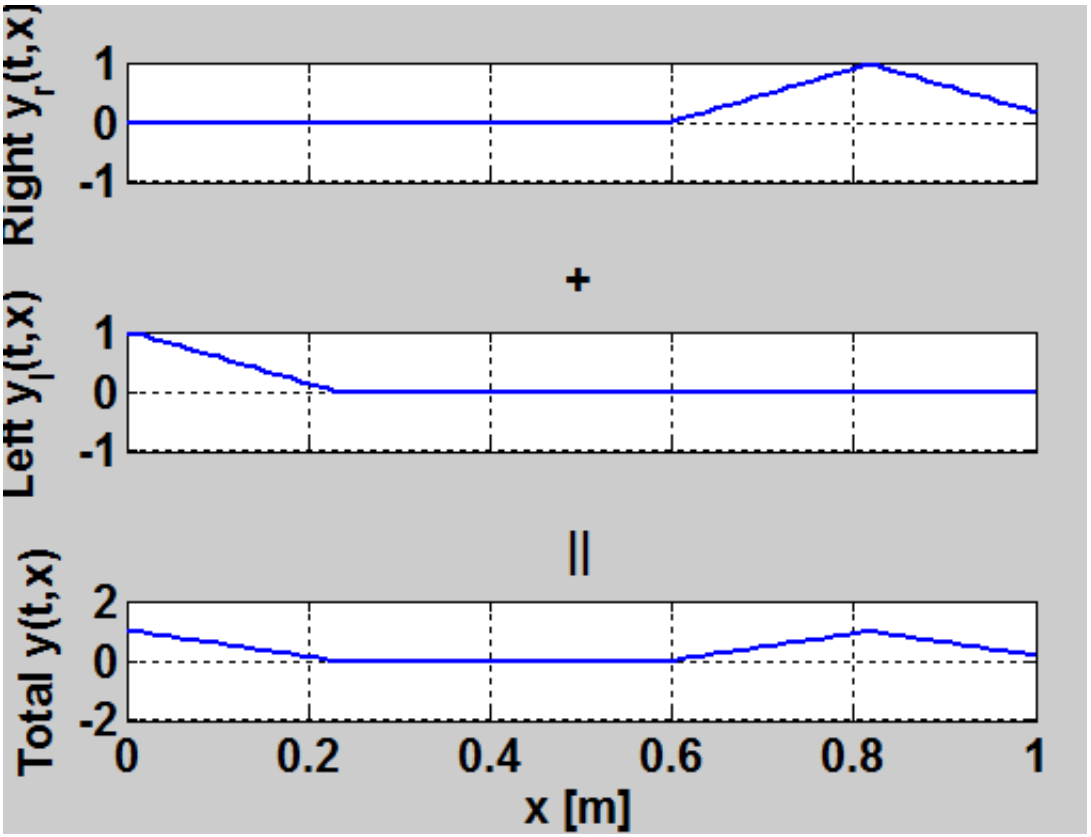

Figure 7. t=20

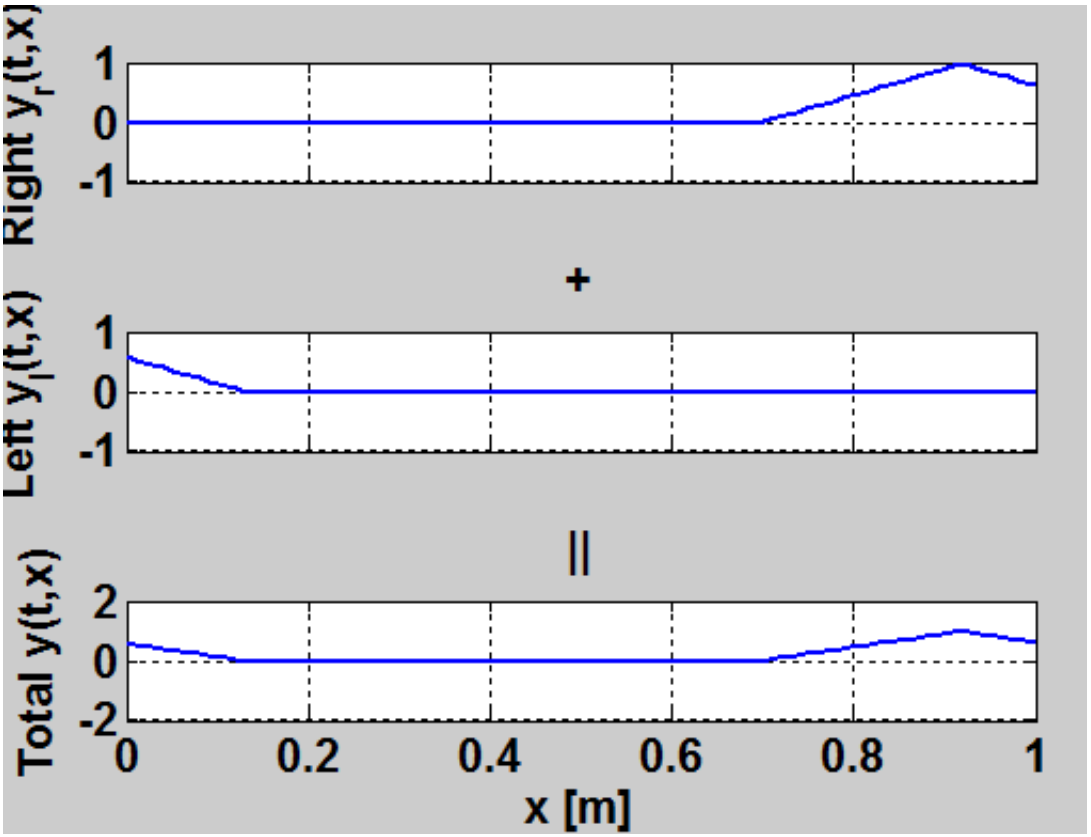

Figure 8. t=25

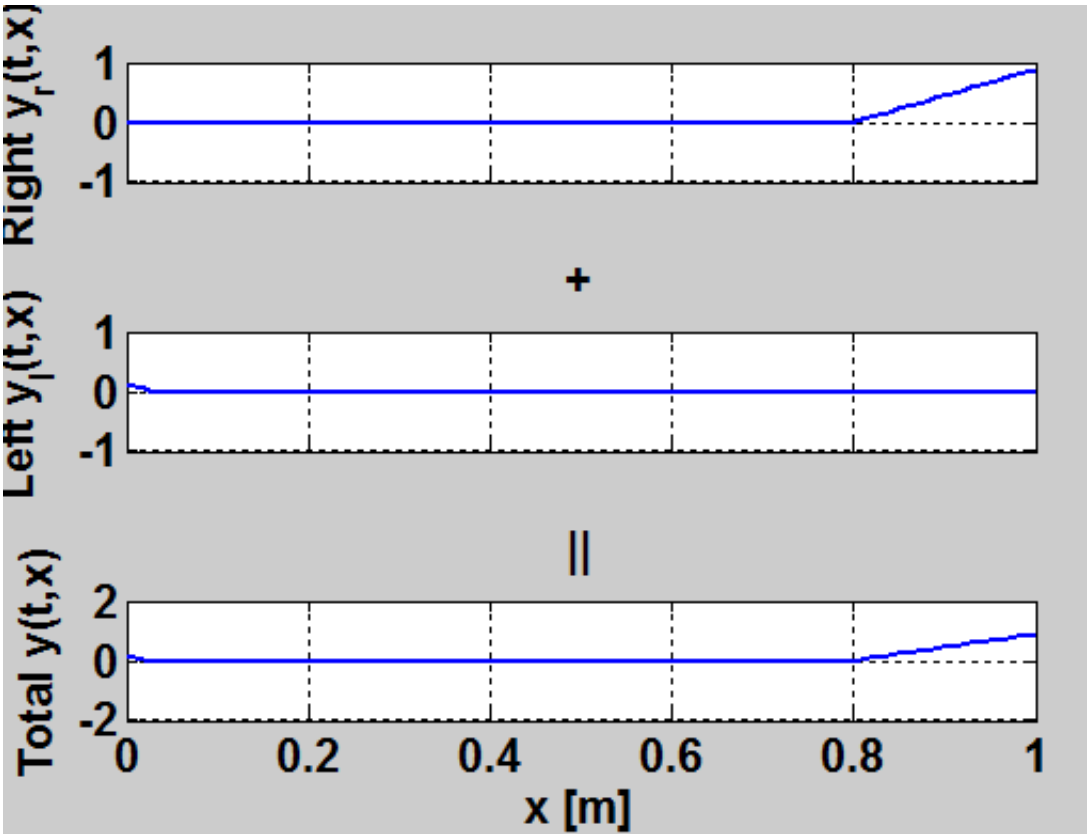

Figure 9. t=30

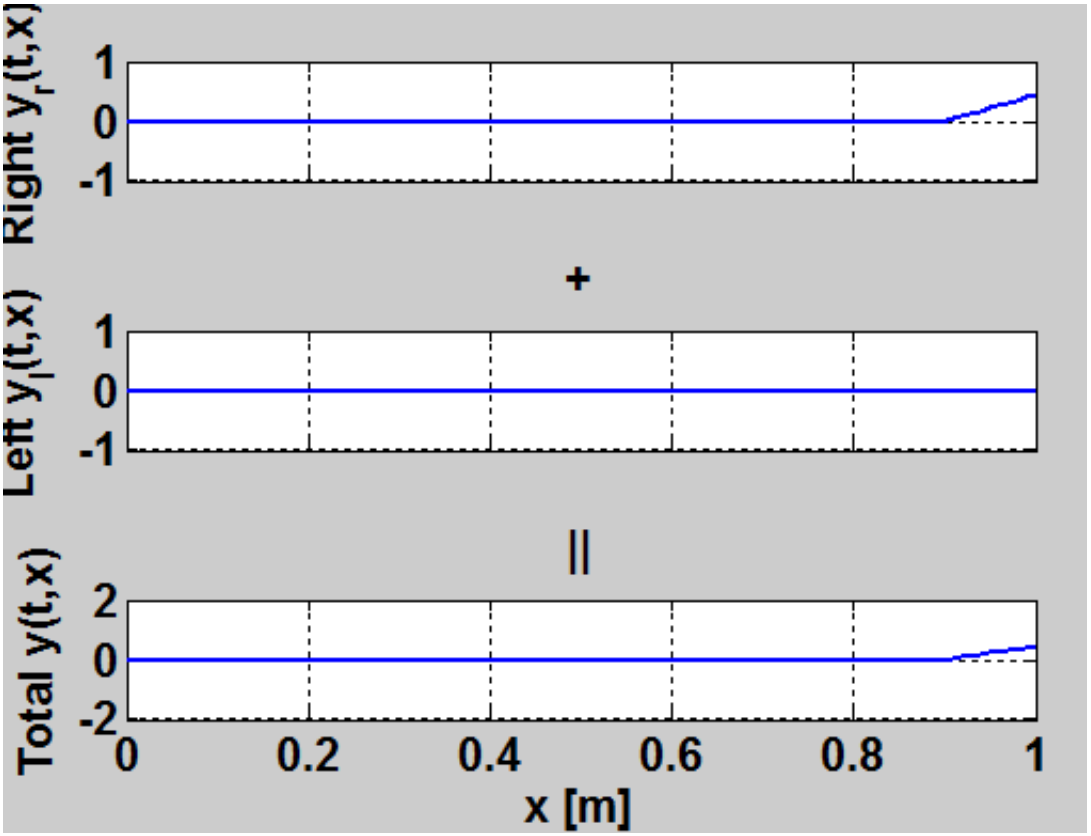

Figure 10. t=35

of simulating propagation distance. Figure 11 shows an example of a length 6 delay line, where at each time slice, values are shifted right. The input signal, into the delay line is $x(t)$, such that the values input into the system are the first six non-zero numbers of the Fibonacci sequence.

As shown in Figure 11, the vector of values is used to represent space. At each tick, or change in time, the values shift right representing propagation of values. For the remainder of the reader, we assume that a delay line represents space and changes occur to its values through time. For simplification, we draw a delay line as a box with the negative exponent to the z being the length of the delay. The z corresponds to the variable of the z-transform, where a single unit delay is z^{-1} [2, 3]. We denote this equivalence in Figure 12.

For the musical acoustics community, such a data structure is invaluable in modeling propagation. In reverberation, delay lines can be used to simulate the indirect paths of a signal [4, 5]. Two online labs are available to help the reader gain familiarity with delay lines and how they are used in modeling propagation distance in a basic feed-forward block diagram and in more elaborate virtual reverberation algorithms. They can be found at http://ccrma.stanford.edu/realsimple/lattice_ladder/ and http://ccrma.stanford.edu/realsimple/reverb/, respectively.

2.1.3. Digital Waveguide Models

Observing the solution to the Wave Equation in Equation 2, the displacement of a string is the sum of two traveling waves moving in opposite directions. We model each traveling wave's propagation path with a delay line. As Figure 13 shows, the data-structure of two delay lines in series in a loop models the physical behavior of the vibrating string as the virtual implementation of the theoretical solution of the sum of two traveling waves.

To cement these abstractions, we present an example of how a vibrating string is modeled with the digital waveguide with initial conditions of a triangular pulse near the middle of the string. Since the sum of the left and right traveling-wave components always equals the displacement of the physical string, we initialize each delay line with a triangular pulse with half the amplitude of the physical string such that the sum of corresponding values from both delay lines equals the amplitude of the pluck for the physical string.

At $t = 0$, we allow values to propagate. Depending on what the sampling-rate is, the physical length of each sample in the delay lines is set assuming a constant propagation speed c.

As Figures 3 to 10 show, the string is held in place at points $P1$, $P2$ and $P3$. At $t = 0$, we release the string at these points. What we obtain from our virtual model is two traveling waves moving in opposite directions such that the sum of

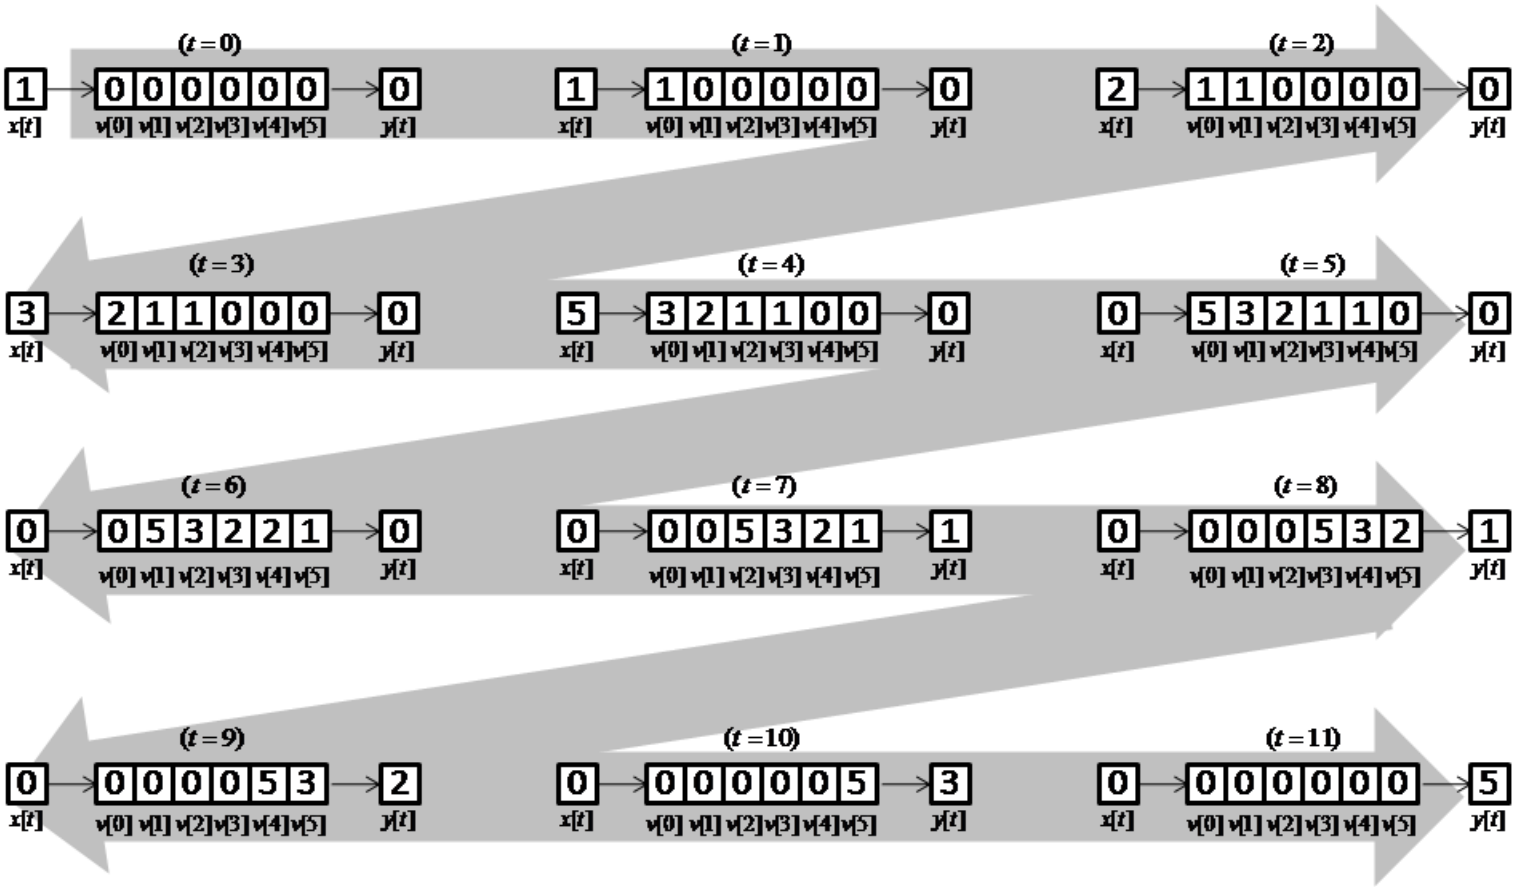

$$x = [1,1,2,3,5,0,0,0,0,0,0,0]$$
$$y = [0,0,0,0,0,0,0,1,1,2,3,5]$$

Figure 11. An example of a delay line and how it is used.

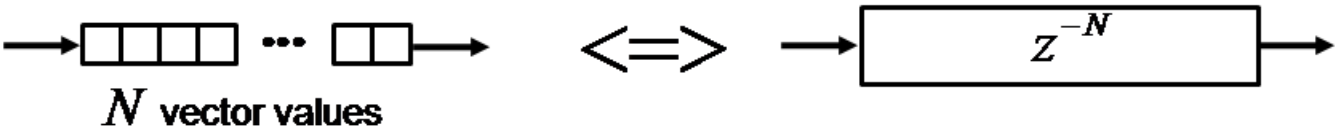

Figure 12. Equivalence between a vector of values and a delay line block diagram.

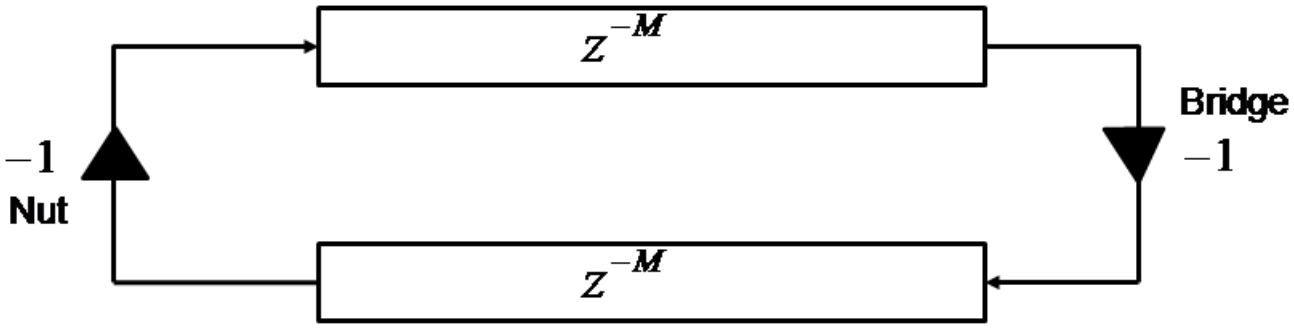

Figure 13. A digital waveguide consisting of two delay lines in series in a loop. Note that the inverters between the delay lines are only present when acceleration, velocity and displacement of the string is modeled.

corresponding values always equals the displacement of the physical string [6].

We now relate the parameters of the digital waveguide to physical characteristics of the string. Our simple digital waveguide, shown in Figure 13, consists of two delay lines each of length M. The frequency of the digital waveguide is directly related to the sampling-rate of our system and the length of the delay lines, M.

Assuming a sampling rate f_s, we can determine how long it would take for an impulse to travel to both ends of the string and back. If an impulse is placed in the first sample of the first delay line, we confirm that it takes $2M$ samples of propagation for the impulse to return to its original position, which is equivalent to $T = \frac{2M}{f_s}$ seconds. Therefore, in our lossless string-model, a traveling-wave takes T seconds to reflect from one end and back. Inverting T, we now know that the virtual string oscillates at a fundamental of $f_1 = \frac{1}{T}$ Hz.

Since each delay line corresponds to one direction of a traveling wave of the string, we have an intuitive physical mapping between the digital waveguide and the physical string. In Figure 13, we have inverters between the two delay lines. The inverter on the left of the figure corresponds to the "nut" of the instrument and the second inverter the "bridge" of the instrument. Note that the inverters are only present in our digital waveguide when acceleration, velocity and displacement waves are modeled with the waveguide. For a more detailed discussion on alternative wave variables modeled in digital waveguides, we refer readers to http://ccrma.stanford.edu/ jos/pasp/Alternative_Wave_Variables.html.

2.1.4. Natural Decay of the String

If a physical string is tied to two rigid ends, energy injected into the string remains within the string. However, it does eventually stop motion. The tension of the string, friction with the air and various effects cause the string's motion to eventually stop [7]. By inserting a low-pass filter into the digital waveguide, as shown in Figure 14, we can model the string's decay properties in our virtual model [8]. Computing the loop filter's coefficients is discussed in Section 4.2.2.

Extending our simple digital waveguide, as shown in Figure 14 we add the string's damping into the Waveguide's feedback loop. This maps physically to energy loss in the string from friction, the string's properties and other more minute, time-invariant effects.

We use a frequency all-pass filter $H_L(\omega)$ to calibrate our Waveguide to match the decay as measured from a real guitar string. In the context of a digital waveguide, $H_L(\omega)$ is called the loop filter.

2.1.5. Modeling Two Planes of Vibration

In the physical world, the strings of an instrument vibrate in two orthogonal planes. In previous sections, we discussed modeling vibrations of the string with a single digital waveguide. To a first approximation, the single digital waveguide representing the transverse motion of the string is accurate, as most energy transferred to the body of the instrument occurs in a single plane. However, to capture more peculiar phenomena, it is necessary to use two digital waveguides, each representing two orthogonal planes of motion. With two digital waveguides and a full covariance-matrix correlating both waveguides, effects such as Two-Stage Decay and beating effects can be correctly modeled [9]. Figure 15 shows the block diagram of two linearly coupled digital waveguides. One digital waveguide, denoted by the vertical direction corresponds to the direction of motion transverse to the instrument's body. The second digital waveguide, denoted by the horizontal direction corresponds to the direction of motion parallel to the instrument's body. Therefore, as a simplification, we often ignore the direct effects of the bridge on the horizontal plane of motion by not having $H_{hh}(\omega)$ in Figure 15.

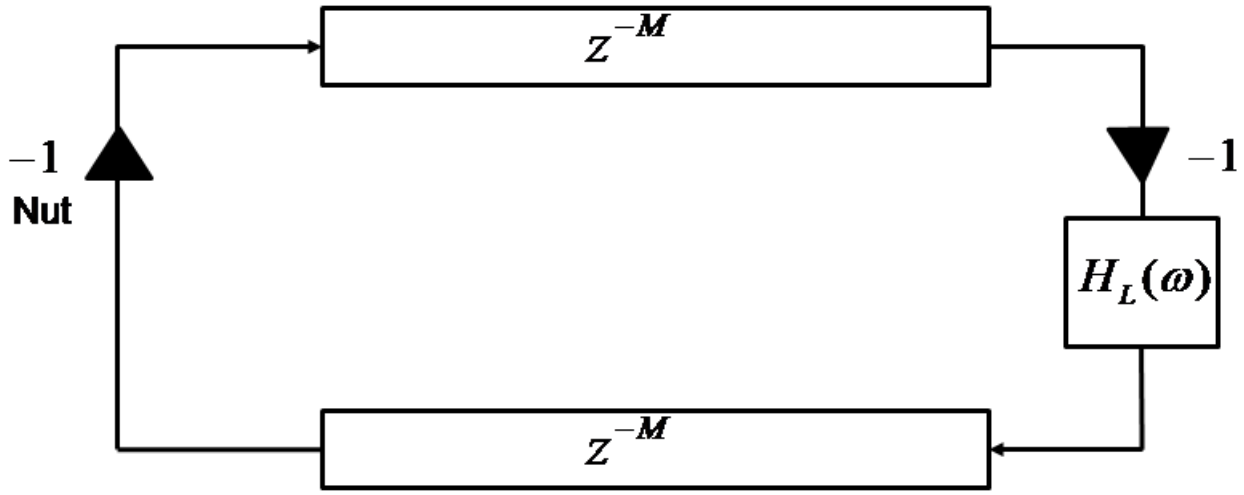

Figure 14. An extension of our simple digital waveguide by including loss through a loop filter. Note that the inverters between the delay lines are only present when acceleration, velocity and displacement of the string is modeled.

Two-stage decay is clearly present in many stringed instruments. Though excitation is mechanically controlled in the piano, two-stage decay is also observed as the strings for a single note are carefully tuned to exhibit what is deemed psychoacoustically pleasing for the instrument. With two digital waveguides, each having different loop-filters for each digital waveguide, one can be set to decay at a faster rate than the other, such that the two-stage decay phenomenon is modeled. In Section 4.2.2, we discuss how to measure the decay of the string from different recordings and tune loop-filters to match the two different decay-rates.

Another phenomenon observed in physical instruments is initial pitch shifts that differ for each plane of vibration [10]. This phenomenon can be modeled using two digital waveguides per string with different time-varying lengths as discussed in Section 2.1.6.

2.1.6. Varying the Digital Waveguide

In Section 2.1.3, we discussed how digital waveguides model propagation of traveling waves along a physical string. The single parameter that can change is the total delay in the loop. Since the described system is linear and time-invariant [3], the digital waveguide's delay lines, presented in Section 2.1.3, can be lumped into one delay line with delay $N = 2M$. Figure 16 shows such a simplification. Though this simplifies the block-diagram of a digital waveguide, it loses its direct physical mapping: viewing the delay-lines as traveling waves moving in opposite directions. However, the systems are equivalent. For our discussion on time-varying lengths for digital waveguides, we discuss digital waveguide's with a single delay line, and thus with a single delay length.

Once reduced to this form, the digital waveguide has a single parameter, its total delay N. By changing N through time, the digital waveguide effectively models a changing propagation distance. The most notable physical everyday occurrence of such an effect occurs whenever an ambulance or police car with its sirens on passes us. The pitch of the siren increases on approach and decreases upon passing. This phenomenon, known as the Doppler Effect [11], is the foundation behind many effects and sounds in the musical community such as the flanger and the Leslie [12, 13]. In the digital domain, however, changes in the delay line length occurs in natural numbers. Fractions of samples in computers do not exist. Therefore, in order to account for varying delay line lengths, many interpolation algorithms have been studied to provide a smooth change in delay line lengths according to the sampling rate used [14, 15, 16, 17, 18]. For a detailed overview of various interpolation methods, we refer the reader to http://ccrma.stanford.edu/ jos/pasp/Delay_Line_Interpolation.html.

An online lab that steps the reader through digital implementations of the Doppler, Flanger and Leslie can be found at http://ccrma.stanford.edu/realsimple/doppler_flanger_leslie/.

2.2. Plucking the String

Returning to our two delay line implementation of the digital waveguide, we show how to excite our virtual string model with a clear mapping to the physical world. Figure 17 shows a digital waveguide where for each delay line, we have tapped into the Xth sample. Here we input at each time-slice our excitation signal $e(n)$.

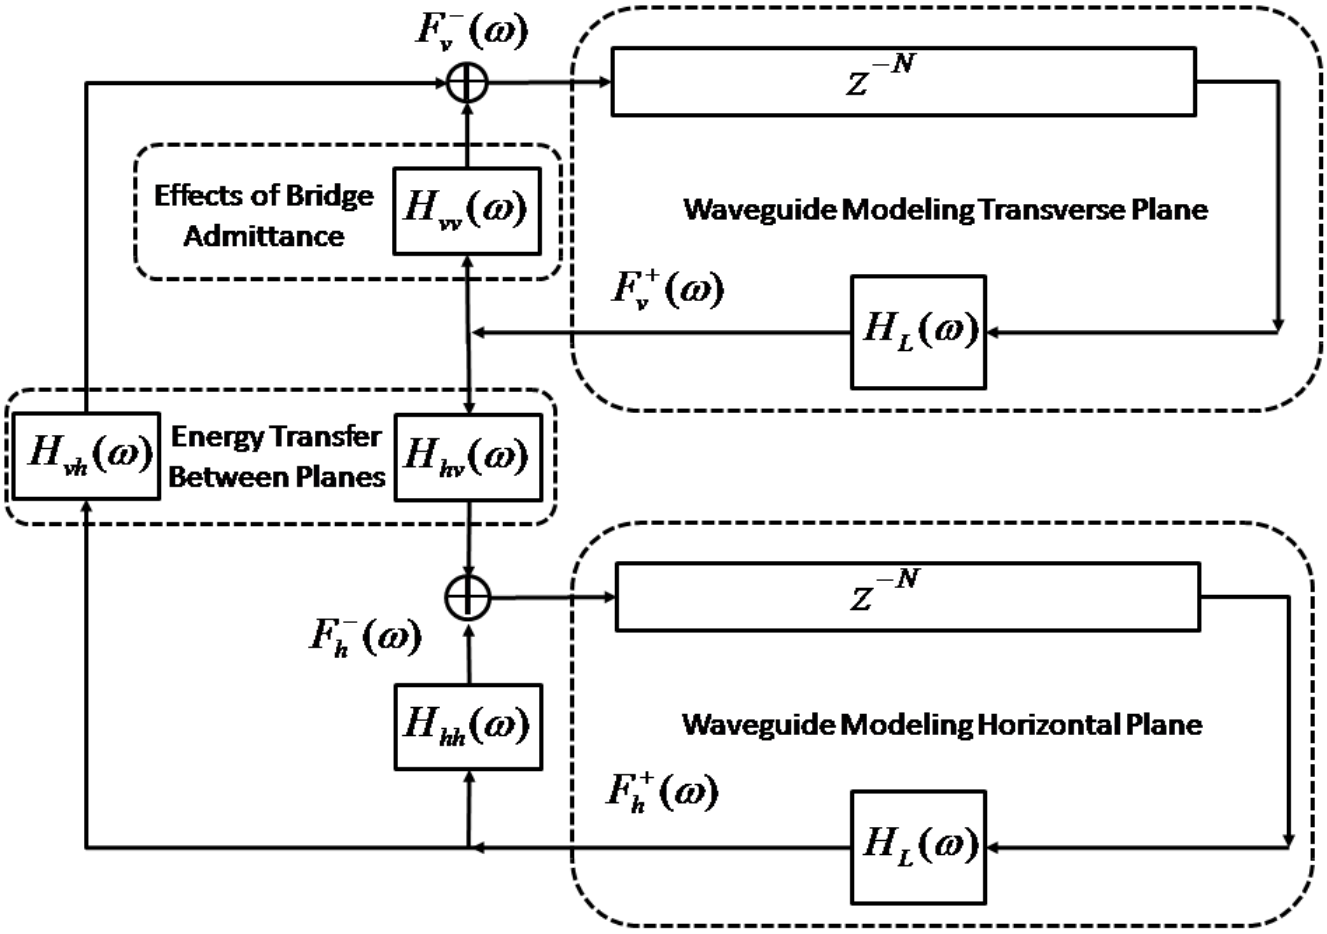

Figure 15. Two linearly coupled digital waveguides. Each Waveguide represents one plane of motion of the string. The vertical plane denotes the transverse direction with respect to the body of the instrument. H_{vv} corresponds to the strings interaction with the instrument's body. H_{vh} and H_{hv} correspond to the coupling between the string in its two orthogonal planes of motion.

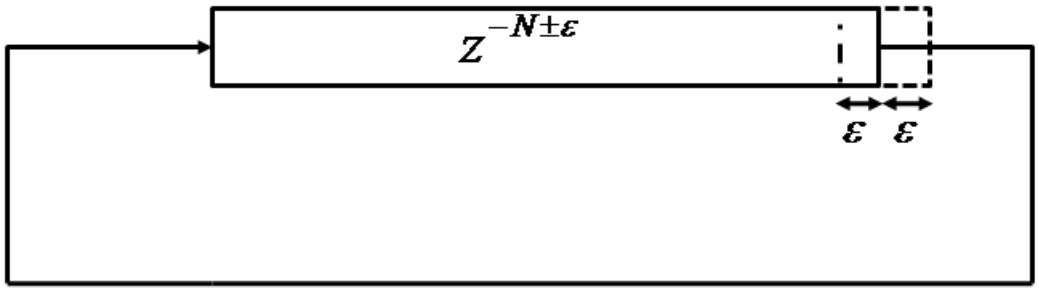

Figure 16. Combining the delay lines in the Digital Waveguide and showing a fraction sample.

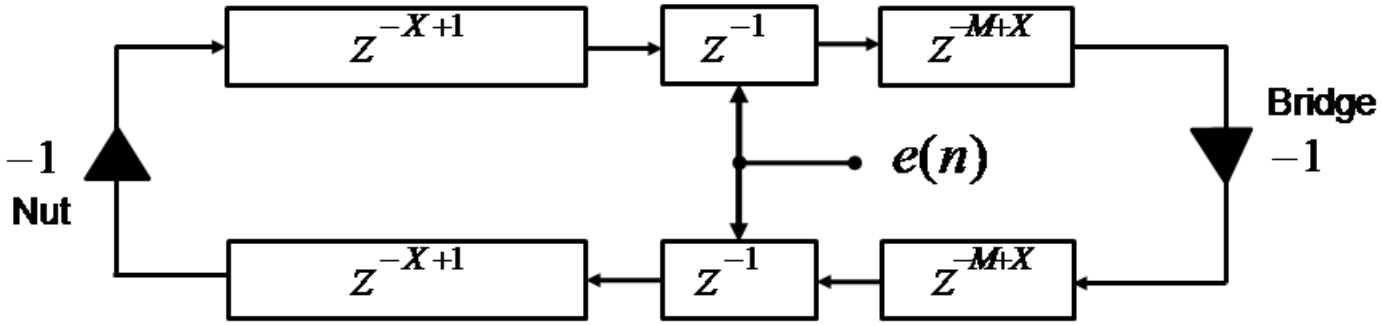

Figure 17. Tapping into the Xth sample where the excitation signal is fed into the digital waveguide.

2.2.1. Theoretical Plucks

If we want an impulse at the sampling distance of where the pluck position of the guitar would be, we calculate X such that the samples for each delay line tapped into correspond to the same physical location of the instrument [19]. We then set $e(0) = 1$ and $e(n) = 0$ for all $n > 0$.

As Karplus-Strong discussed, exciting a digital waveguide with an impulse is psychoacoustically plain, whereas exciting the digital waveguide with random values results in a more satisfying sound [20].

With delay lines modeling displacement, we can linearly-interpolate between the pluck-point to the two ends of the string for both traveling waves for an excitation that displaces the string at a single point causing the remainder of the string to displace linearly to the ends of the string. Furthermore, if the delay lines represent acceleration or curvature, this ideal-pluck is a single non-zero sample in each delay line. Note, since the digital waveguide is a sampled model of a continuous system, band-limiting conditions need be accounted for. A detailed description of such can be found in http://ccrma.stanford.edu/ jos/pasp/Ideal_Plucked_String.html.

2.2.2. Complexities of Real Plucks

For theoretical plucks, exponentially-decaying short-noise-bursts produce the psychoacoustically richest sounding digital waveguides for plucked stringed instruments. However, as discussed later in Section 4.1, the pluck itself is psychoacoustically significant. As any plucked-stringed instrument musician would attest to, reproducing the exact same pluck by hand is extremely difficult if not impossible. Studies have been made to study guitar plucks [21]. In such research, elaborate devices have been constructed to reproduce the exact same pluck. As shown later in Section 4.1, plucks are more than short bursts of white-noise. Each pluck has unique characteristics to its particular occurrence.

With coverage of idealized plucks and digital waveguides, one can now build an electric guitar using digital waveguides. An online laboratory outlining the steps for doing so can be found at
http://ccrma.stanford.edu/realsimple/electric_guitar/.

2.3. Body Resonance

This section describes the integration of the instrument body into our digital waveguide model. We extend our digital waveguide string model with the instrument's driving-point admittance. With the inclusion of the instrument body's resonator, we can attach multiple instantiations of our digital waveguide model to account for multiple strings coupled at the bridge of the instrument [22].

2.3.1. Driving-Point Admittance

The driving-point admittance of the body of the instrument, characterizes linear-interactions between the string and the body in both directions: meaning energy transferred from both string to body and body to string. The driving-point admittance is defined as follows:

$$\Gamma(\omega) = \frac{V(\omega)}{F(\omega)} \tag{3}$$

where for a given frequency ω, the ratio between the Fourier Transform of the velocity and force is known. A physical explanation of the driving-point admittance is to view the admittance as a measure of how readily force exerted at the contact point at a certain frequency results in motion at that same frequency [23]. Since the body of the guitar is a

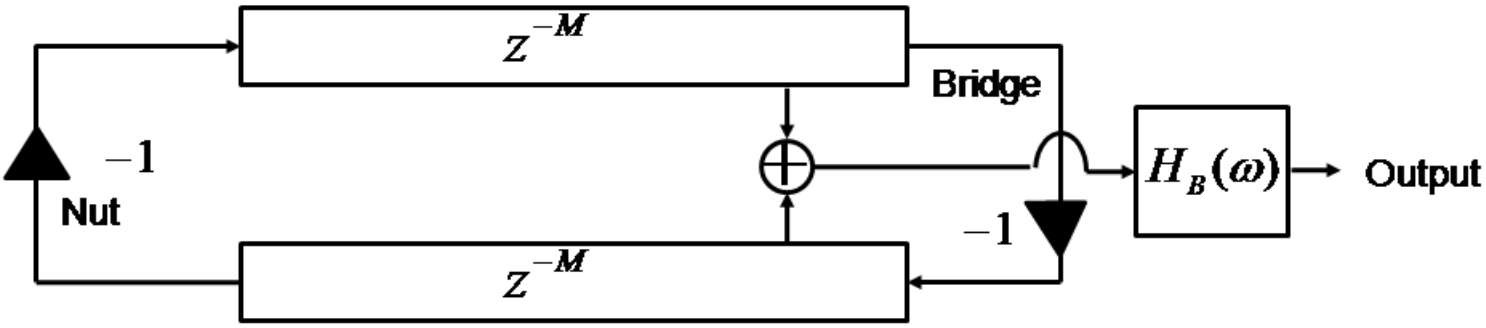

Figure 18. Convolving the physical output of the digital waveguide with the driving-point admittance.

rigid structure that exhibits standing waves at particular frequencies, known as the modes of the guitar, the body vibrates naturally at these frequencies. Notable modes include the air mode and the first few body modes.

A brief overview of the mechanics of our physical model thus far: the string vibrates upon excitation. Its end is connected to the body of the instrument at the bridge, exerting energy from the initial pluck at the fundamental frequency and its harmonic series. According to the driving-point admittance, the force applied at the bridge by the string results in motion of the bridge and top-plate. The resulting motion is dependent upon the construction of the top-plate of the instrument which determines its modes of vibration. The acoustic instrument then propagates pressure waves according to its top-plate movement, thereby coloring the resulting pressure waves heard by our ears.

2.3.2. Filtering with the Driving-Point Admittance

The simplest extension of our current digital waveguide model is to use the driving-point admittance of the instrument in a convolutory way. A discussion of what measurements needed to compute the driving-point admittance can be found in Section 3.2.

In Figure 18 the output of the digital waveguide is convolved with the driving-point admittance. Since the virtual instrument described up to this point is linear and time-invariant, the driving-point admittance of the instrument can be used to excite the digital waveguide by commuting the admittance with the digital waveguide. Though this is unrealizable in the physical world, the two systems are the same. This method of using the driving-point admittance to drive the digital waveguide is known as *commuted synthesis*. [1] [2]

2.3.3. Bi-directional Interaction

In the physical world, the motion of the string exerts a force on the bridge resulting in a velocity at the point of contact. In Section 2.3.2 from a physical perspective, only energy passing from the string to the body is modeled. However, we can extend our current model to account for not only the influence of the string to the bridge, but also the influence of the bridge to the string. Network theory applied to this junction results in a bi-directional model [24, 25, 26]. Figure 19 is an illustration of six strings attached to a common bridge. The simplifying assumption is that all strings move at the same velocity as the bridge at the bridge.

$$H_B(\omega) = \frac{2}{R_b(s) + R_1 + R_2 + ... + R_6}$$

Where $R_b(s) = \frac{1}{\Gamma_b(s)}$, where $\Gamma_b(s)$ is the driving-point admittance of the bridge and R_i is the wave impedance of string i. Furthermore, each string's wave impedance can be computed using its physical characteristics through the following equation:

$$R_i = \sqrt{K_i \varepsilon_i} = \frac{K_i}{c_i} = \varepsilon_i c_i$$

where K_i is the string tension, ε_i is the string's linear mass density and c_i is the speed at which waves on the string travel both left and right. For more information on Network theory and how it applies to musical acoustics, we refer readers to http://ccrma.stanford.edu/ jos/pasp/Introduction_Lumped_Models.html.

[1] http://ccrma.stanford.edu/~jos/pasp/Commuted_Synthesis_Strings.html

[2] An online laboratory for creating an acoustic guitar physical-model using commuted synthesis is available in the RealSimple collection. [3]

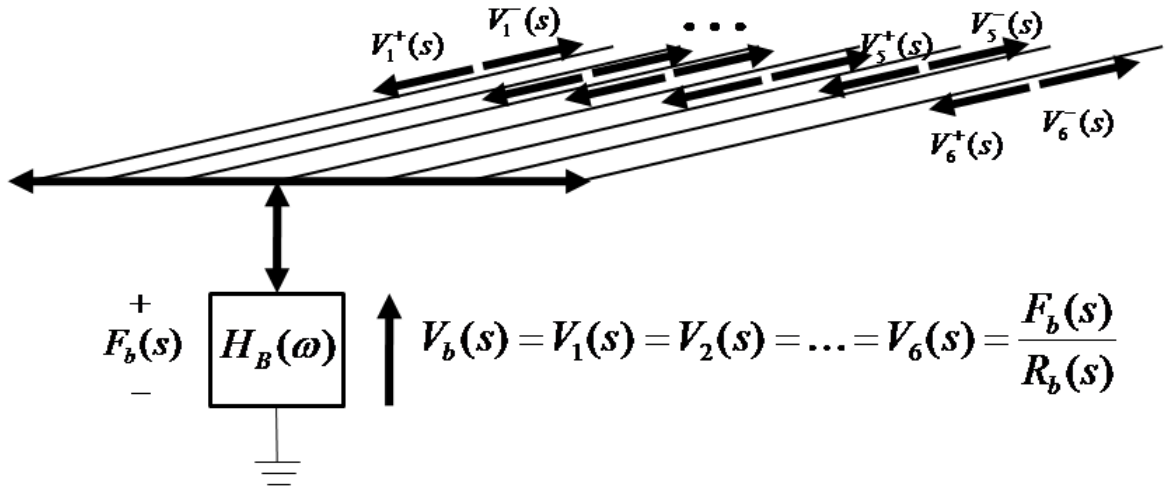

Figure 19. An illustration of six strings attached at the bridge of the instrument. The simplifying assumption is that all strings move at the same velocity of the bridge at the bridge.

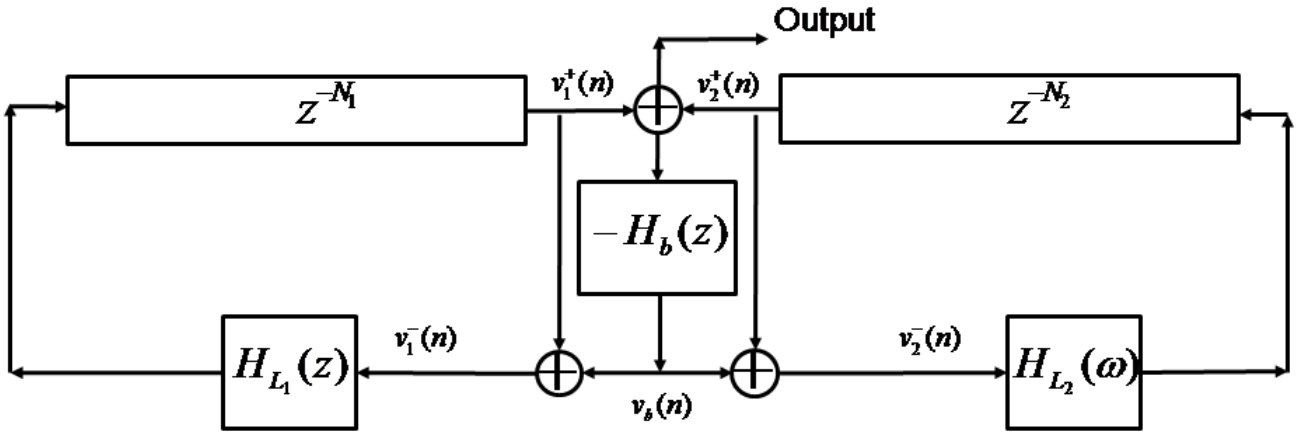

Figure 20. A block diagram of multiple strings of an instrument connected at a single scattering junction with the body of the instrument.

2.3.4. String-Body Scattering Junction

Continuing our discussion in Section 2.3.3, to model multi-stringed instruments, we naturally extend our model so that each digital waveguide, models a string of the instrument, with all strings attached to a scattering junction with a load bearing at the junction, representing the mass of the instrument's bridge and top-plate. Figure 19 shows a diagram of our virtual model of a multi-string instrument with all strings coupled at the bridge.

With such a model, we are able to physically model certain interactions occurring at the bridge between the strings and body of the instrument. In the case of the acoustic guitar, all the strings are coupled at the bridge such that energy transfers from one string to another through the bridge. This effect is very prominent, in that musicians change their technique to account for "leakage" of energy. Guitar players refer to this phenomenon as open-string noise, which typically occurs when a guitar player plays on higher strings without muting lower strings, resulting in lower strings vibrating without being plucked. This is especially problematic for musicians playing microphoned acoustic instruments which often result in unwanted feedback. For electric guitar players, when the gain of the pickups are amplified to create distortion, unwanted vibrating strings obscure the sounds of the actual strings played.

Figure 20 shows a block-diagram modeling an instrument with two strings with a string-body scattering junction. The digital waveguides representing each string are simplified in that they move only in the transverse plane. As shown, energy from one string can cause excitation of another string [27]. Furthermore, the resulting output contains the effect of the

body with the driving-point admittance acting as a mass load at the junction. The digital waveguides used in Figure 20 can be easily replaced with the digital waveguide in Figure 15 such that the output of $H_{vv}(\omega)$ from each digital waveguides' transverse loop feeds into a summer into $-H_b(z)$ of Figure 20. The output of $-H_b(z)$ is then fed back into each digital waveguide's transverse plane loop [28].

2.4. Pressure Radiation

This section describes two radiation models for the guitar. In the first approach, the sound pressure is synthesized at any number of virtual microphone positions by means of a simplified Kirchhoff-Helmholtz integral over the radiating surfaces determined by the modal state [29]. The second approach is a simplified point-to-point method which uses recorded input/output signals to compute a transfer function between the string-bridge junction to the position of the output-recording microphone.

2.4.1. Integrating over Radiating Surface Areas

Soon to come!

2.4.2. Point-to-Point Modeling

This method of modeling uses a single transfer function, which can be either computed analytically or measured to compute what the resulting pressure waves at a given point in space from the point of vibration are from acceleration waves. The limitations of this method, compared with Section 2.4.1 is that pressure waves can be computed for only one point in space, relative to the vibrating source. Where as described in Section 2.4.1, the radiation of sound is computed for the source's surrounding space. Though this method measures radiation for the source's surrounding space, it comes with a higher computation cost. In real-time synthesis, point-to-point modeling is the preferred method as computing the entire radiating surfaces of a source is too costly to be done in real-time. Furthermore, for a stationary binaural listener, only two transfer functions for a stereo signal need be computed.

3. MEASUREMENTS

This section reviews techniques for guitar measurement relevant to building a virtual guitar.

3.1. String Vibration

String vibration measurement has been studied and used extensively in the music industry in the form of amplifying guitars, resulting in the electric guitar. In the traditional single-coil magnetic pickup, a strong magnet is wound with copper wire. This acts as a sensor as the motion of the string causes a change in the magnetic field and as a result a change in the current through the copper wire. This signal, proportional to the movement of the string, is then amplified. For some guitars, such as Gibson's Les Paul, the guitar is close to being purely electric, in that energy is maintained within the vibration of the string to drive the pickups that amplify its sound. The guitar body is literally a block of solid wood meant to not resonate or color the timbre of the guitar's sound, thereby minimizing energy tranferred to the guitar body to create longer sustain in the string.

Another method for measuring a string's displacement is with use of a light-emitting-diode (LED) and a photo-transistor. The photo-transistor measures the amount of light received from the LED as light is obstructed by the movement of the string. This sensor is an inexpensive method for measuring the motion of the string in two-orthogonal planes.

Vibrometry is another method used for measuring the motion of a string as it measures the frequency shift of the back-scattered light reflected from the measured surface to compute the changing distance between the laser and the vibrating string at high sampling rates.

3.2. Bridge Force

Measuring the driving-point impedance of a stringed instrument's body is key in obtaining its acoustic characteristics [30, 31]. To obtain the driving-point impedance of the guitar body, an impedance head accelerometer is placed in contact with the guitar on or close to the bridge. This sensor outputs the accleration at its point of contact. Accelerometers are typically very low load bearing thereby obtaining accurate measurements up to very high frequencies.

As defined in Section 2.3.1, the velocity and force are used in computing the driving-point-impedance. In the following subsection, we present two methods for exciting the motion of the guitar bridge. With each method, the force signal is measured and used in conjunction with the acceleration measured to compute the driving-point impedance of the bridge.

3.2.1. *Excitation Force*

We present two methods for exciting the bridge for measurements. One uses a shaker and the other an impulse-force hammer.

The shaker is placed in contact with the bridge with as little load mass from the shaker. A signal, typically white noise is input to the shaker. Both the force signal from the shaker and accelerometer signals are recorded to compute the driving-point admittance.

With the impulse-force hammer, an individual strikes the bridge with a short, single instance to excite the bridge. The impulse-force hammer outputs the force signal applied to the testing body, and the accelerometer signal measures the movement of the bridge. With both signals, similar to the shaker and accelerometer method, the driving-point admittance is computed.

There are benefits and disadvantages to both excitation methods. Since transfer function computation is statistical in nature, in that the transfer function is computed using the cross power spectral density, and since the signal used to drive the shaker is white noise, the measured signals can be recorded indefinitely. Therefore, the resulting transfer function is less susceptible to experimental noise. However, the drawback to the shaker is its hefty mass relative to the impulse-force hammer. Obtaining an accurate response at higher frequencies is more difficult with a shaker.

The benefit of the force-impulse hammer is that the mass of the hammer is not as much of an issue as the shaker in exciting the body of the instrument. Therefore, the high-frequency components of the resulting transfer function are more accurate. However, the drawback with the force-hammer is in its difficulty in reproducing the exact same measurement. It is difficult if not impossible to strike the body of the instrument with the same force by hand. Furthermore, the resulting signals are short compared to what can be obtained when a shaker is used.

3.3. Body Vibration

Understanding the movement of the main resonating surface of an instrument, in the case of a guitar, its top-plate, is essential in understanding the instrument's ability to propagate sound at all frequencies. Viewing the frequencies at which the modes occur helps explain the timbre of the acoustical characteristics of the instrument. A scanning-point vibrometer, for example by Polytec, measures the movement of a surface and gives its users three-dimensional software views of its structural motion.

3.4. Pressure Radiation

To measure a point-to-point transfer function for pressure radiation, the instrument is taken into an anechoic chamber, and excited with a shaker with white noise as input or with a force-hammer, in which case the force input into the system is close to an impulse. The signal generated from the motion of the instrument is recorded using a microphone a distance away from the instrument. With the acceleration and pressure signals recorded, a transfer function is computed [32].

4. PARAMETER ESTIMATION

This section reviews methods for calibrating the model elements discussed in Section 2 to match measurements discussed in Section 3.

4.1. Excitation

As discussed in Section 2.2, the excitation to our virtual instrument is a signal, taken from a wave-table used to drive our digital waveguide models. In the literature there are several methods for computing excitation signals from actual recordings. Extracting excitations from actual recordings of guitar plucks result in the best, psychoacoustically sounding models [33].

In recent years, two fundamentally different approaches for obtaining excitation signals from recordings have emerged: inverse-filtering and constant overlap-add (COLA) methods. We briefly describe the problem and give an overview of the various methods.

Figures 21 and 22 show a recorded guitar note on the guitar's high 'e' string and its Short-Time-Fourier-Transform (STFT), respectively. As shown, upon the onset of the note, all frequencies have energy. After the inital attack, most of the energy remains at the harmonic frequencies of the fundamental. The goal of excitation extraction is to remove the tonal components that ring after the initial onset and to reduce the energy during the onset at those components to match the general energy levels at other frequencies.

With the problem now graphically represented, we describe how the two different methods approach removal of harmonic peaks shown in the STFT. The inverse-filtering methods remove the harmonic peaks by inverse-filtering the signal

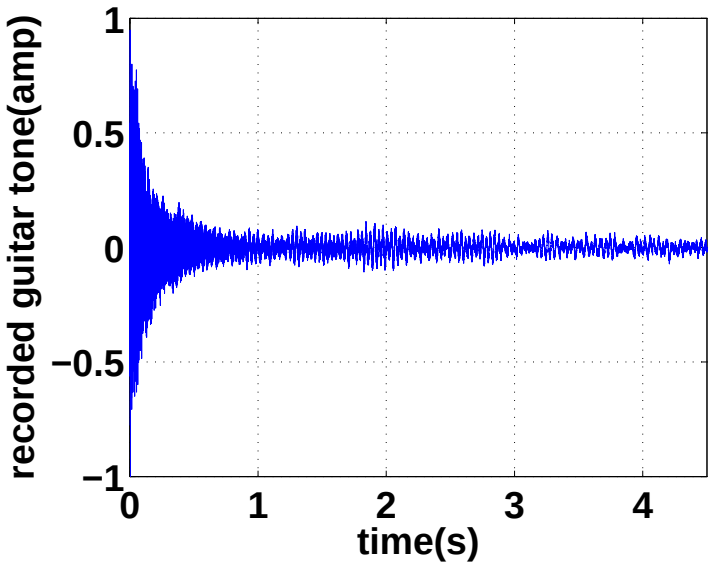

Figure 21. Time-domain plot of a recorded guitar tone on the open high ’e’ string.

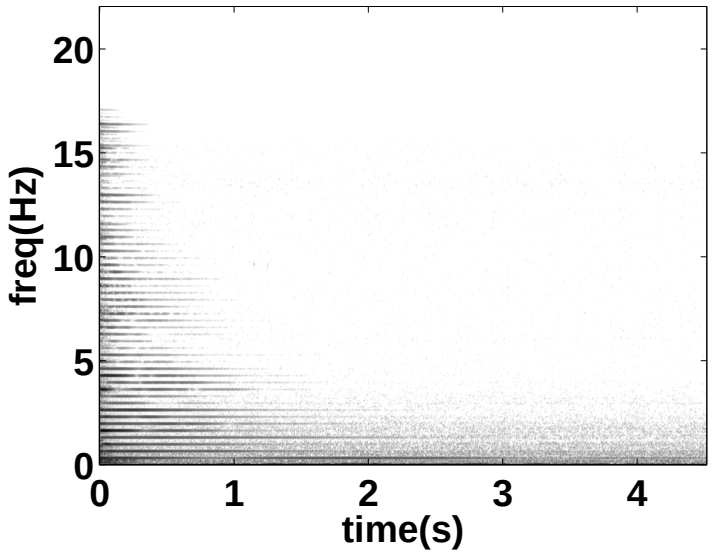

Figure 22. Spectrogram of a recorded guitar tone on the open high ’e’ string.

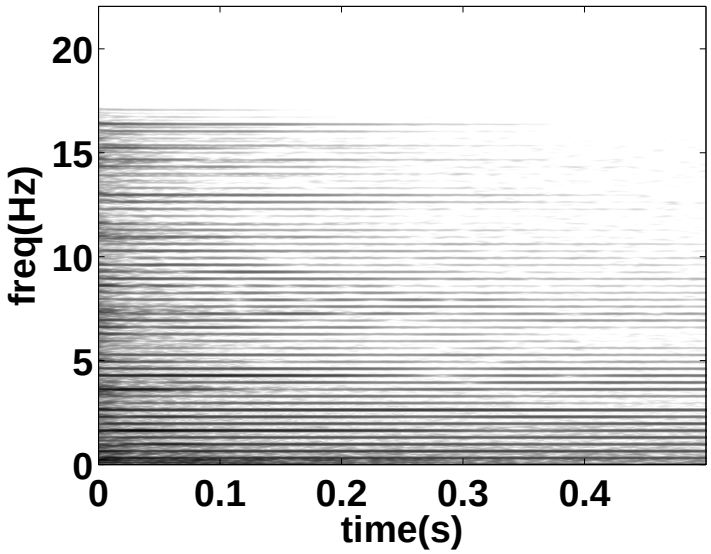

Figure 23. Spectrogram of a recorded guitar tone on the open high 'e' string from time $t = 0$s to time $t = 0.5$s

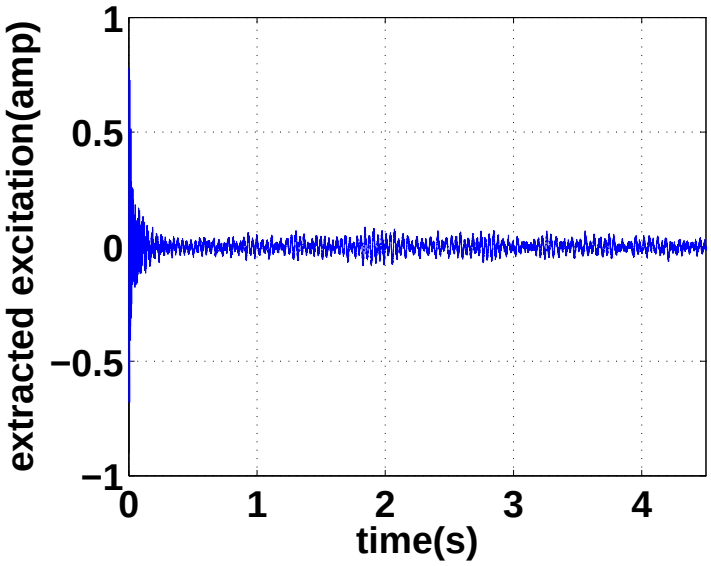

Figure 24. Time-domain plot of the same recorded guitar tone with harmonic peaks removed resulting in an excitation signal for physical model use.

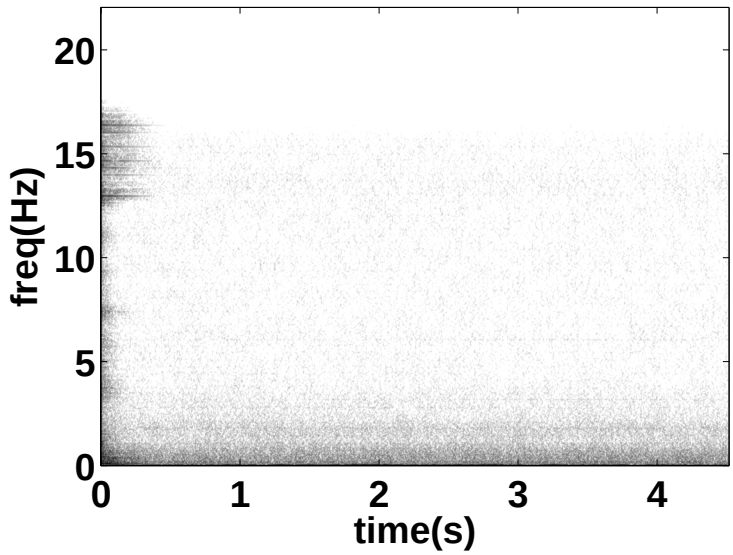

Figure 25. Spectrogram of the same recorded guitar tone with harmonic peaks removed resulting in an excitation signal for physical model use.

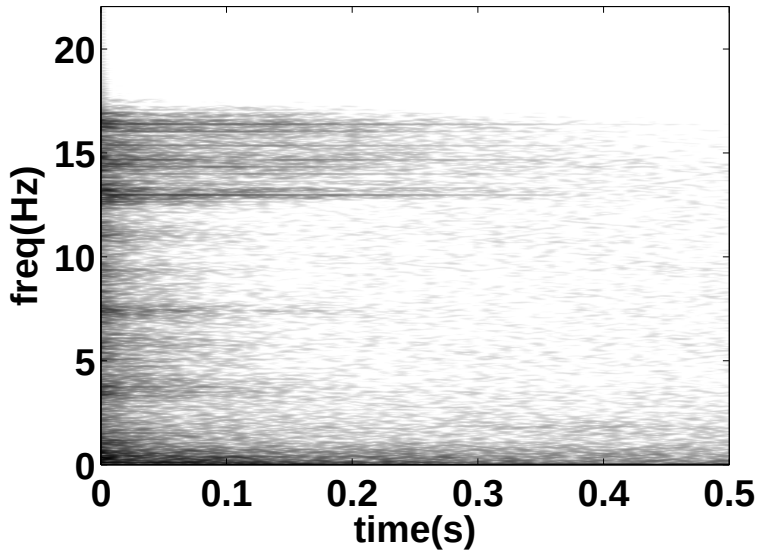

Figure 26. Spectrogram of the same recorded guitar tone with harmonic peaks removed resulting in an excitation signal for physical model use from time $t = 0$s to time $t = 0.5$s.

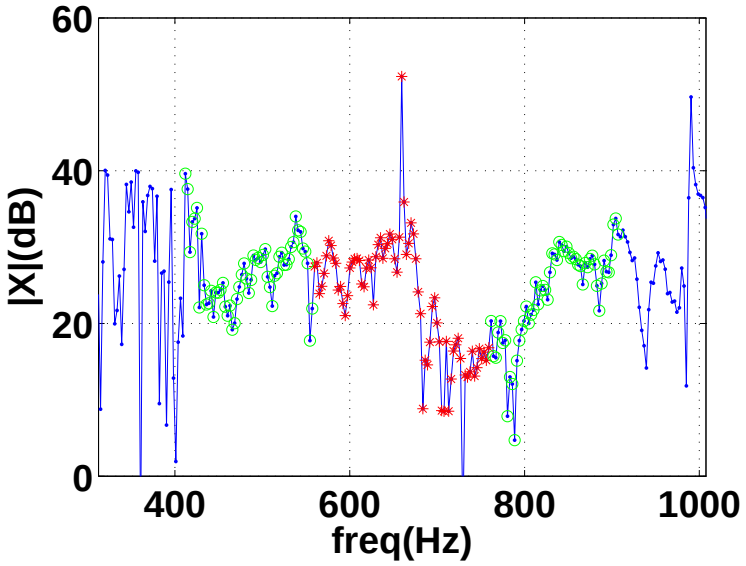

Figure 27. Plot of the original recorded guitar tone's first FFT frame with focus on the peak near 660Hz. Circle dots are FFT values to be used for statistics collecting.

with a comb-filter with peaks at the harmonic frequencies. The COLA methods remove the harmonic peaks by applying non-linear averaging to the magnitudes at each STFT frame of the signal.

The first method is the Matrix-Pencil Inverse-Filtering method [34]. It computes the sinusoidal components of a signal, using the Matrix-Pencil method, and performs inverse-filtering with the sinusoidal components to remove the tonal components leaving the excitation [35].

The second method, the Sines-Plus-Noise Inverse-Filtering method is similar to the Matrix-Pencil method, in that sinusoidal components of a model are computed, but instead of using the Matrix-Pencil method for computing the sinusoidal components, a generative sinusoidal model is used [36, 37]. A residual signal is computed from subtracting from the original recording the sinusoidal signal. Inverse-filtering is then performed on the recording using a Digital Waveguide tuned for the recording where scaled versions of the residual and sinusoidal signals are added together to help remedy the notches created from inverse-filtering [38].

The third method, the Magnitude Spectrum Smoothing, uses STFT processing. Within each FFT window, a low-pass filter is applied to the magnitude spectrum of the window. The iFFT is then taken where the resulting time-signal is stored in a buffer for overlap-add [39].

The last method, the Statistical Spectral Interpolation method, similar to the MSS method, performs STFT processing and COLA reconstruction, but removes harmonic peaks by sampling new spectral magnitudes at the peaks according to a normal distribution with mean and covariance equal to those of magnitudes in the values surrounding the peak for each FFT frame [33].

As the literature shows, SSI produces the best psycho-acoustically sounding excitations. Here, we present the method in detail.

From a high-level viewpoint, the SSI method only modifies the magnitudes of the STFT of the guitar tone without affecting phase information. The method collects statistics on the magnitudes of frequencies surrounding harmonic peaks and uses these statistics to generate non-deterministic gain-changes for the magnitudes at these peaks, without modifying the phase. With inverse-filtering methods, modifying phase inevitably introduces artifacts. Thus, this method's primary goal is to minimally-alter the original tone.

The STFT is used for analyzing and modifying the original recorded tone. The STFT can be seen as a sliding window that takes at each sample-window a FFT of the windowed signal. The transform of that windowed portion is then modified, and the iFFT is then taken and saved in a buffer. The window is then slid according to how much overlap is wanted. The parameters for the STFT are the type of window used, the length of the window and the number of samples the window slides by. Sample parameters for the method are a Hamming window of length 2^{12} samples with 0.9 overlap (hop size of 410 samples). Though 2^{12} samples with a sampling-rate of $44,100$ Hz is long (≈ 100 ms), pre-echo-distortion artifacts are remedied by starting the algorithm during the onset of the recorded tone.

Actual processing occurs at each window of the STFT. Consider each FFT window taken to be a frame for processing. Within each frame, the harmonic peaks of the recorded tone are attenuated. Harmonic peaks can be found using the Quadratically Interpolated FFT (QIFFT) method [40].

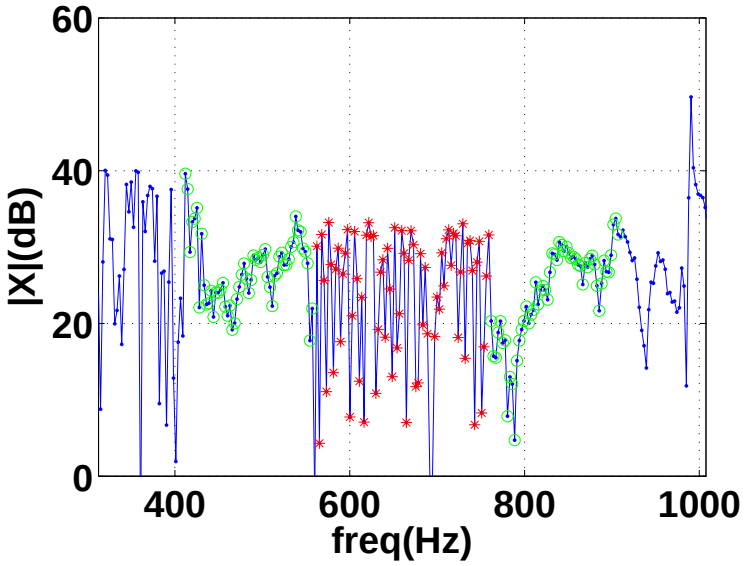

Figure 28. Plot of the original recorded guitar tone's first FFT frame with focus on the removed peak near 660Hz. Circle dots are FFT values to be used for statistics collecting.

Assuming that the fundamental frequency of the recorded tone is at f_1 in Hz. A bandwidth W_p in Hz is specified, indicating the width of the peak. Another bandwidth W_n in Hz is specified indicating the width of the interval used for statistics collecting with respect to the fundamental frequency f_1. In using the SSI method for the recording in Figure 21, $W_p = 0.3 \cdot f_1$ and $W_n = 0.75 \cdot f_1$. These values ensure that the points used for statistics collecting do not reach into the next harmonic peak but are large enough to obtain a reasonable mean and standard deviation.

For each harmonic i with frequency f_i, the following is defined and used for processing.

Define a set of indices, Γ, whose frequency values satisfy the following:

$$\forall \gamma \in \Gamma, W_p \leq |\nu_\gamma - f_i| \leq W_n \tag{4}$$

where ν_k corresponds to the frequency in Hz of the kth FFT bin. The values in Γ correspond to indices within the current frame whose frequencies lie within the specified band W_n but outside the band W_p centered around f_i. See the circled points in Figures 27 and 28.

The mean and standard deviation of the magnitude of values in FFT bins in Γ are computed as follows:

$$\mu = \frac{1}{|\Gamma|} \sum_{i \in \Gamma} |X_i| \tag{5}$$

$$\sigma = \sqrt{\frac{1}{|\Gamma|} \sum_{i \in \Gamma} (|X_i| - \mu)^2} \tag{6}$$

Define the set of indices, Δ, whose frequency values satisfy the following:

$$\forall \delta \in \Delta, |\nu_\delta - f_i| \leq W_p \tag{7}$$

The values in Δ correspond to indices within the current frame whose frequencies lie within the specified band W_p centered around f_i. The magnitudes at these frequencies are changed to remove the peak. See the starred points in Figures 27 and 28.

Thus, for all bins with indices in Δ, magnitude values are modified to remove the observed peaks. This occurs as follows:

For each $\delta \in \Delta$, generate a value $\rho \sim \mathcal{N}(\mu, \sigma)$.

$$X_\delta := \frac{\rho}{|X_\delta|} X_\delta. \tag{8}$$

Figures 27 and 28 shows the points the algorithm uses for statistics collecting and the points with gains altered. As shown, the peak at 660Hz is entirely removed.

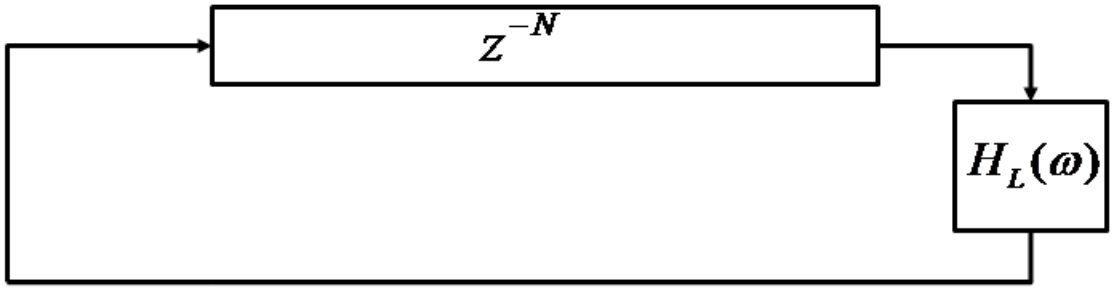

Figure 29. Digital Waveguide model with one single delay line and loop filter.

4.2. String

4.2.1. Single Plane of Vibration

4.2.2. Loop Filter Estimation

As discussed in Section 2.1.4, the loop filter gives the Digital Waveguide frequency-dependent damping. Similar to what is used in Artificial Reverbation, the time for frequencies to decay 60dB, known as t_{60}, is used to determine the gains of the filter. We define t_{60} to be the time in seconds for a signal to decay 60dB. For obtaining gains for the fundamental and harmonics of our loop filter, we compute t_{60} for each of these frequencies over time using the Energy Decay Relief (EDR) [36, 41].

For the remainder of this section, we commute the second delay line shown in Figure 14 with the loop filter and combine the delay lines. The resulting system, shown in Figure 29 is equivalent to the one shown in Figure 14 as the delay line and loop filter components are linear and time-invariant.

Similar to the STFT, the EDR shows the spectral components over time of a given signal. The single difference between the STFT and the EDR is that the square magnitudes of the FFT of each frame is used instead of the magnitudes. With the EDR, the amount of energy in each FFT bin is given over time. An advantage of using the EDR over the STFT is that energy decreases monotonically, whereas magnitudes, as used by the STFT, may increase and decrease resulting in noisier data for decay estimation.

The EDR is defined as follows

$$EDR(t_n, f_k) = \sum_{m=n}^{M} |H(m,k)|^2 \tag{9}$$

where $H(m,k)$ corresponds to the kth FFT bin at frame m of the EDR. M corresponds to the total number of frames in our EDR. Thus, $EDR(t_n, f_k)$ represents the total amount of energy remaining in the signal at time $t = t_n$ for frequency f_k. Figure 30 shows a three-dimensional plot of the EDR over time and frequency.

Since the displacement of the string decreases exponentially, viewing its motion and EDR in log space, we expect linear decays, as Figure 30 shows for a plucked guitar string note.

For each frequency f_k in the set of frequencies closest to the fundamental and harmonic frequencies, we fit a line to the time-decay of its EDR. The primary parameter for estimation is the slope of the line. Figures 31 to 36 show EDR plots for the fundamental frequency and five harmonics. The EDR for each frequency over time is plotted in blue. The estimated line is shown in red. The estimation used only the segments of the EDR for which the red line overlaps the blue. As shown, the line estimated nearly matches exactly the subset of EDR points used.

Given the estimated slope, m_k, for frequency f_k, $t_{60}(k) = \frac{-60}{m_k}$. With $t_{60}(k)$ computed for the fundamental and harmonic frequencies, we now relate the decay times to the gains of the loop filter.

$$|G(\omega_k)|^{\frac{t_{60}(k)}{T}} = 0.001 \tag{10}$$

where $|G(\omega_k)|$ corresponds to the per sample desired gain of our filter. We also have the relation that

$$|H(\omega_k)_L| = |G(\omega_k)|^M \tag{11}$$

$$|H(\omega_k)_L|^{\frac{1}{M}} = |G(\omega_k)| \tag{12}$$

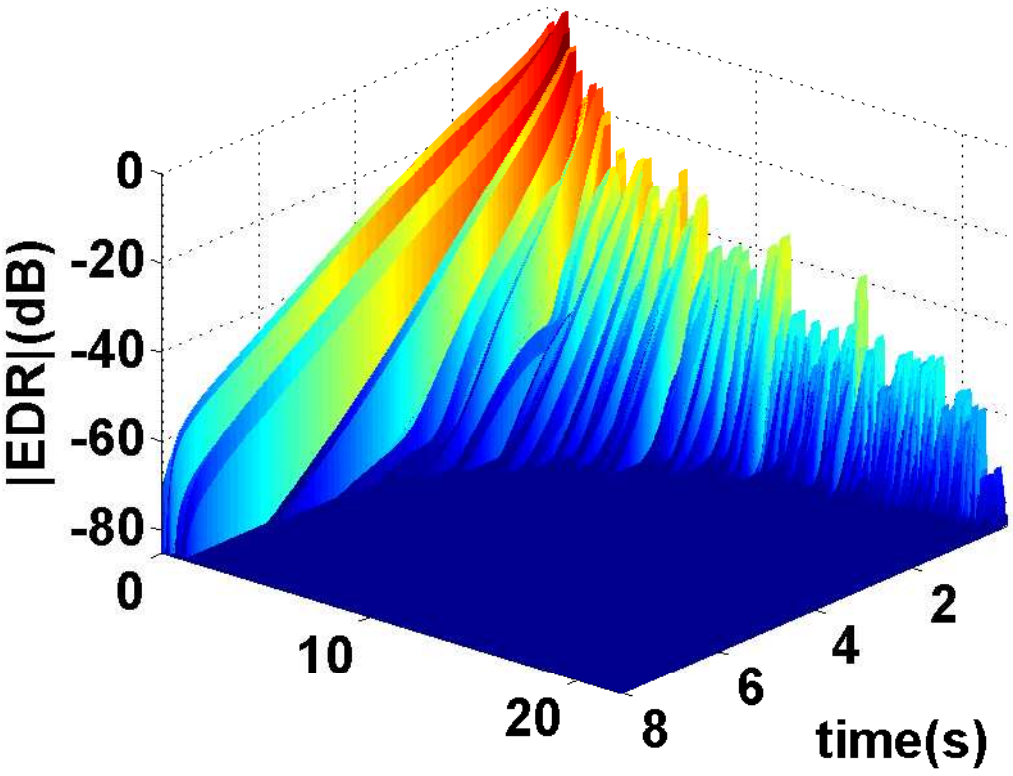

Figure 30. Plot of the Energy Decay Relief of a plucked guitar note.

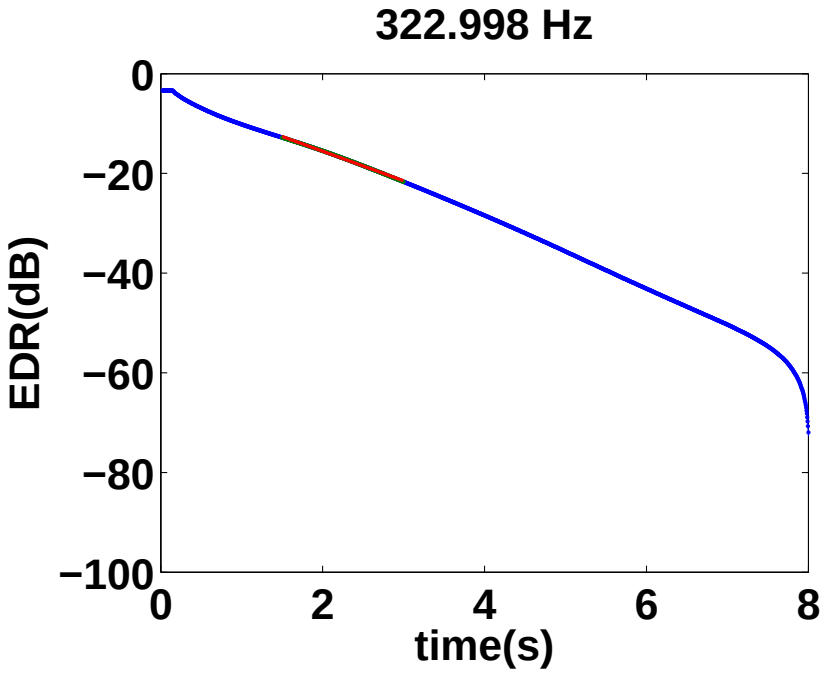

Figure 31. EDR over time at frequency bin 322.998Hz. The overlapping red segments were fitted with the underlying blue points and plotted using the estimated slopes.

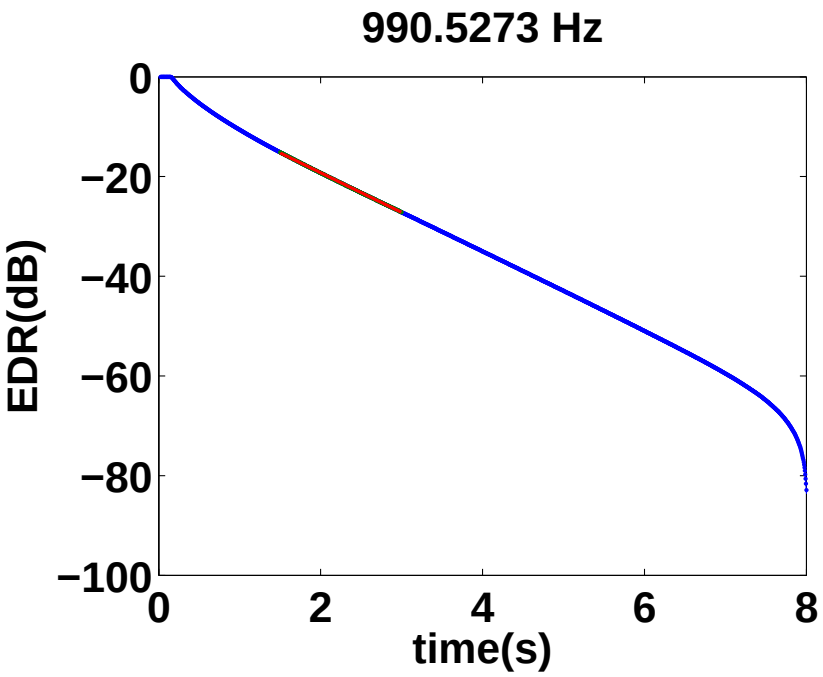

Figure 32. EDR over time at frequency bin 990.5273Hz. The overlapping red segments were fitted with the underlying blue points and plotted using the estimated slopes.

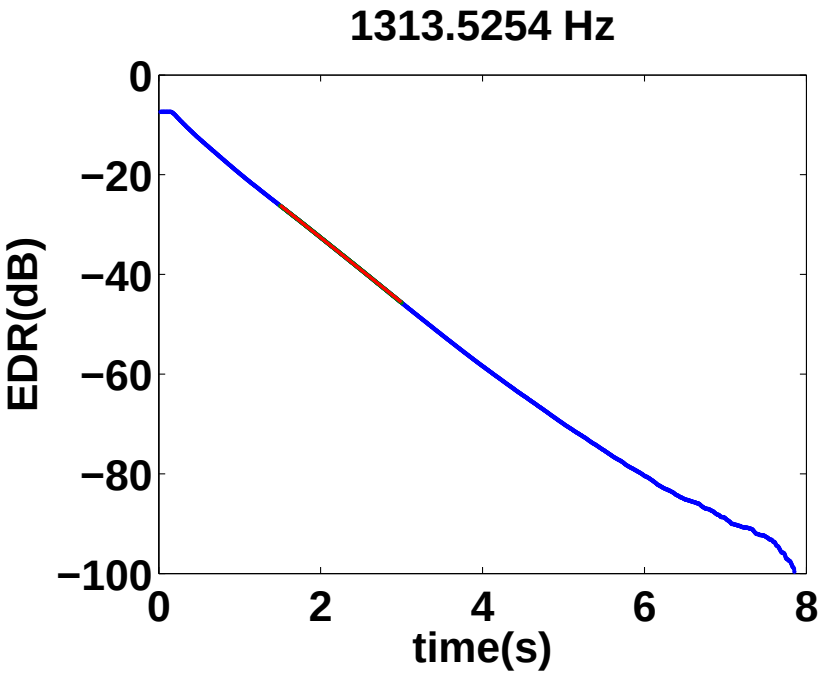

Figure 33. EDR over time at frequency bin 1313.5254Hz. The overlapping red segments were fitted with the underlying blue points and plotted using the estimated slopes.

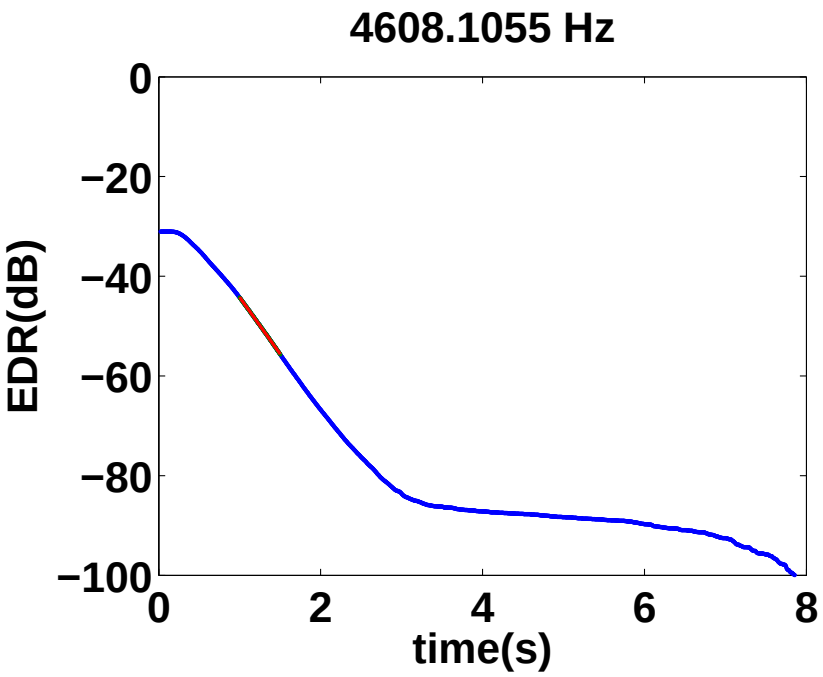

Figure 34. EDR over time at frequency bin 4608.1055Hz. The overlapping red segments were fitted with the underlying blue points and plotted using the estimated slopes.

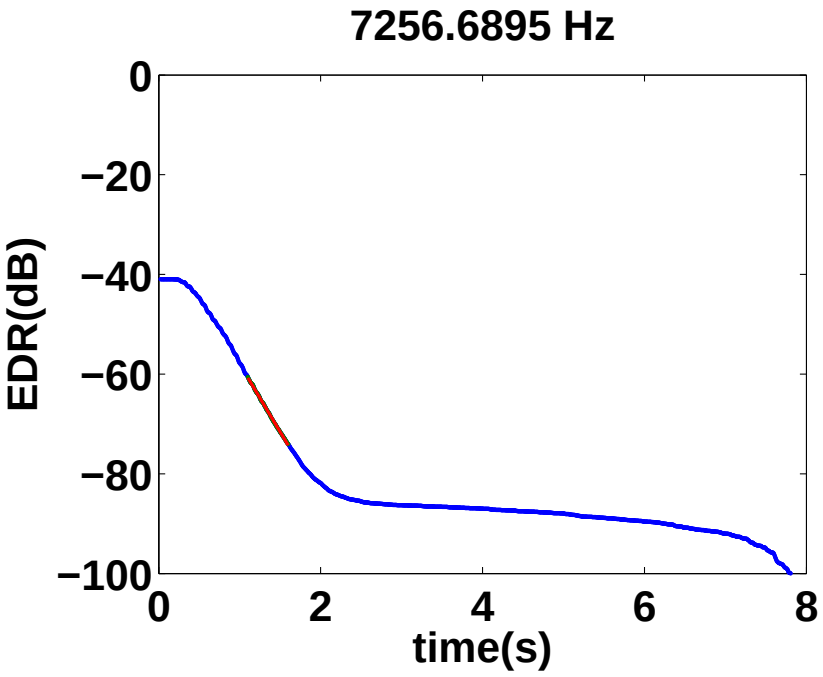

Figure 35. EDR over time at frequency bin 7256.6895Hz. The overlapping red segments were fitted with the underlying blue points and plotted using the estimated slopes.

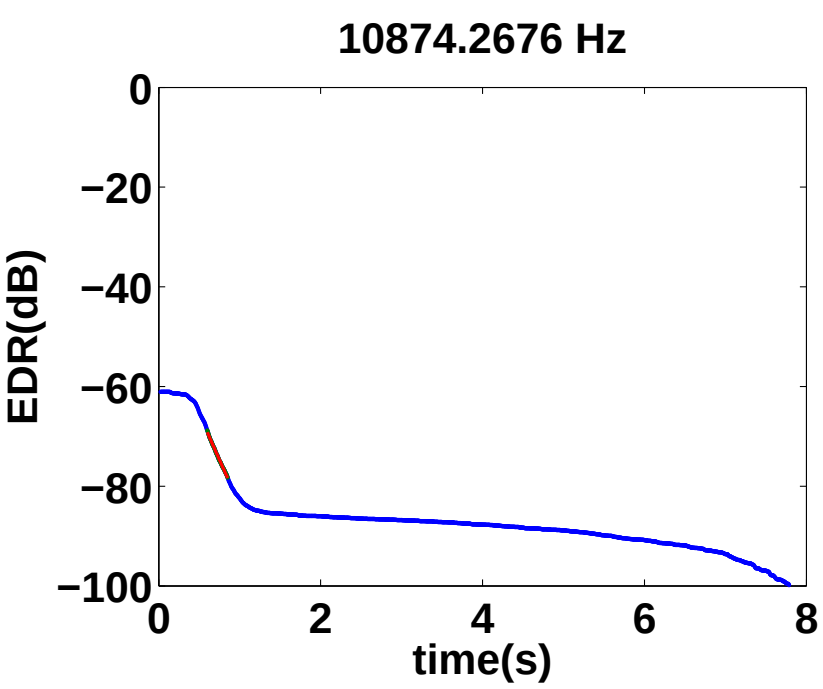

Figure 36. EDR over time at frequency bin 10874.2676Hz. The overlapping red segments were fitted with the underlying blue points and plotted using the estimated slopes.

where $|H(\omega_k)_L|$ is the desired loop filter gain. This relates our per sample gain with the desired loop filter gain in series with a length M delay line. Substituting $|H(\omega_k)_L|$ for $|G(\omega_k)|$ in Equation 10, we obtain

$$|H(\omega_k)|^{\frac{t_{60}(k)}{MT}} = 0.001 \tag{13}$$

Taking $20\log_{10}$ of both sides yields

$$20\log_{10}|H(\omega_k)| = -60\frac{MT}{t_{60}(k)} \tag{14}$$

To specify gains at frequencies between the fundamental and harmonics, we linearly interpolate between points for which we have gains for. In specials cases, the gains at DC up until the fundamental frequency are set to the value of the gain computed at the fundamental. For frequencies between the last harmonic and half the sampling-rate, the gains decrease at an arbitrary slope.

With the gains specified for the loop filter, a complex spectrum with minimum-phase is computed [42]. A least-squares fit or the Steiglitz-McBride iteration is used to compute the coefficients for an arbitrary order filter [42, 43].

Figures 37 and 38 show the desired gains computed for a recorded guitar note. The fundamental is 330Hz. Gains for the fundamental and harmonic frequencies are circled in red. All other points are linearly interpolated between these gains and the constants we set for DC and half the sampling-rate. Gains below the fundamental to DC are set to the value of the gain at the fundamental frequency. The gain at half the sampling-rate was arbitrarily set to -4.08dB. Figure 37 shows gains from DC up to 5kHz.

4.2.3. Two Orthogonal Planes of Vibration

Soon to come!

4.3. Body Resonator

From our synthesis model, the body resonator has a driving-point admittance that filters the output of the string model to transform force waves along the string into acceleration waves at the bridge [44]. Taking the inverse-FFT of the measured driving-point admittance, we obtain the impulse response of the body of the instrument. This time-domain signal can be stored as a wavetable for convolution with the output of the synthesis string-model. Figures 39, 40, 41 and 42 show the time-domain signal and spectrum (both magnitude and phase) of a measured guitar's body impulse response, respectively. As shown, the impulse response lasts well over 50ms long. Viewing its spectrum in Figure 41, we note peaks that occur near 120 and 250Hz.

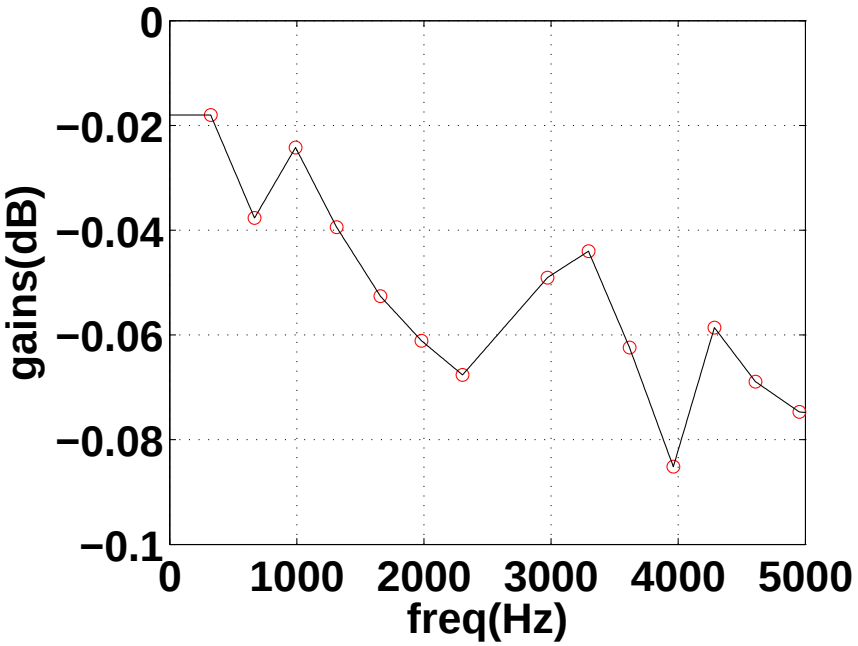

Figure 37. Estimated gains from a guitar's high 'e' string. Only frequencies below 5kHz are shown. The fundamental frequency is 330Hz. Circled points correspond to gain values for the fundamental and harmonic frequencies. All other points were linearly interpolated. The gain at DC was set to be the same as that of the fundamental. The gain at half the sampling-rate was set arbitrarily to -4.08dB.

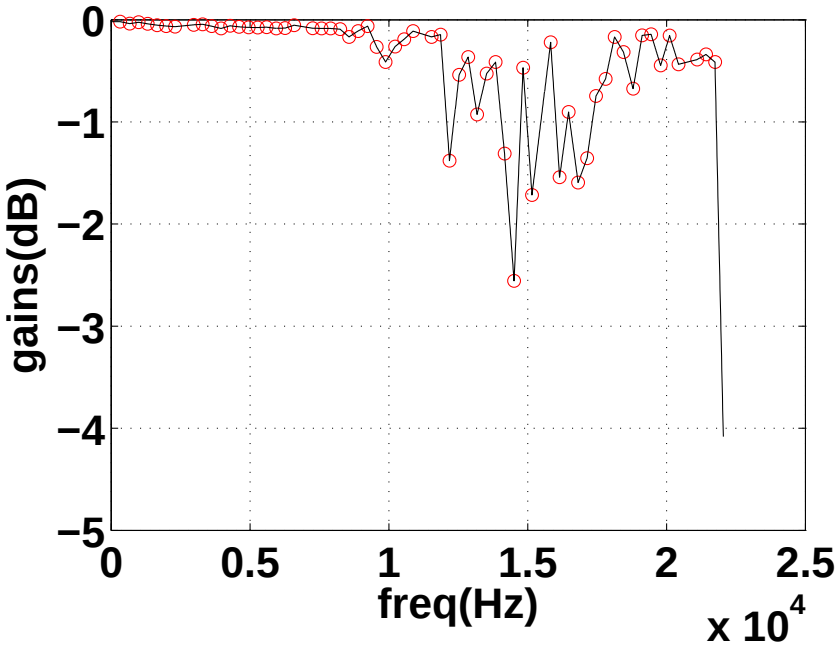

Figure 38. Estimated gains from a guitar's high 'e' string. The fundamental frequency is 330Hz. Circled points correspond to gain values for the fundamental and harmonic frequencies. All other points were linearly interpolated. The gain at DC was set to be the same as that of the fundamental. The gain at half the sampling-rate was set arbitrarily to -4.08dB.

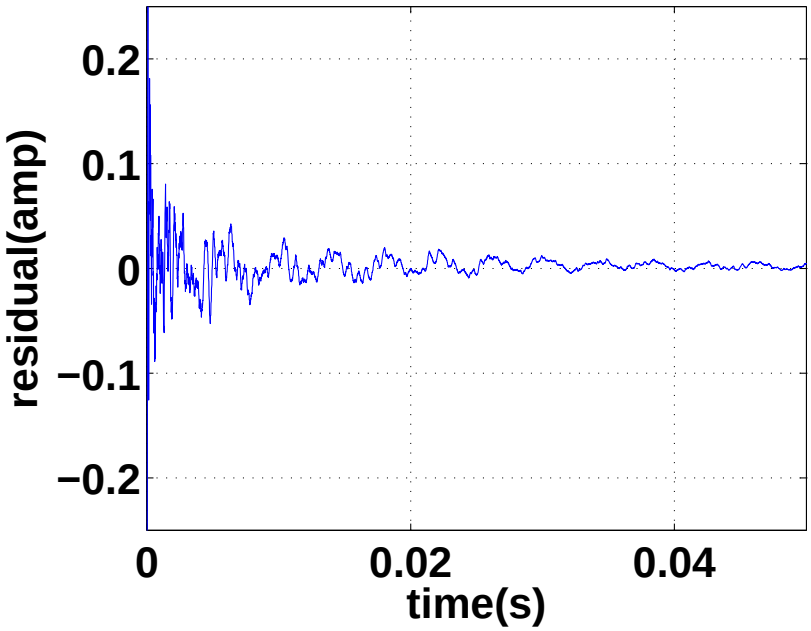

Figure 39. Measured body response from a gypsy guitar. The plot shows the signal from time 0 to 50ms.

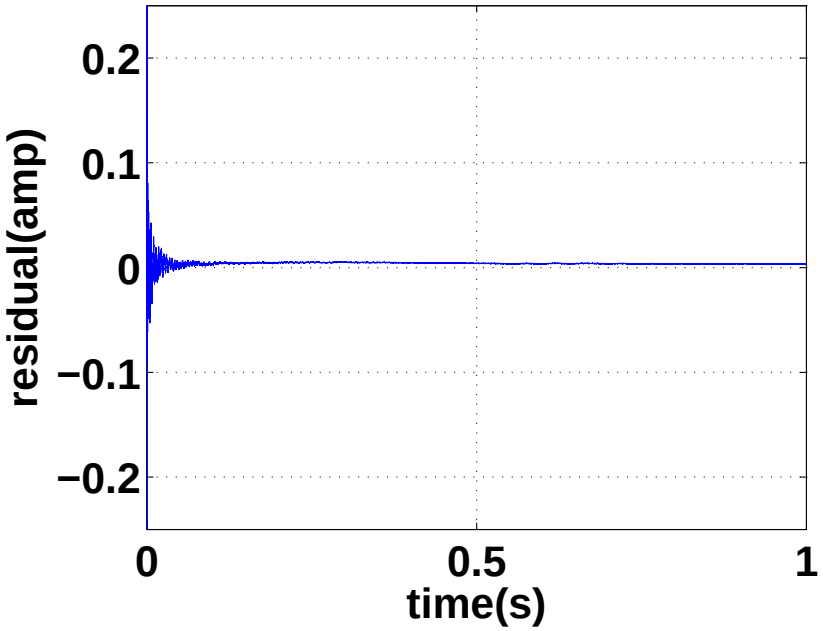

Figure 40. Measured body response from a gypsy guitar. The plot shows the signal from time 0 to 1s.

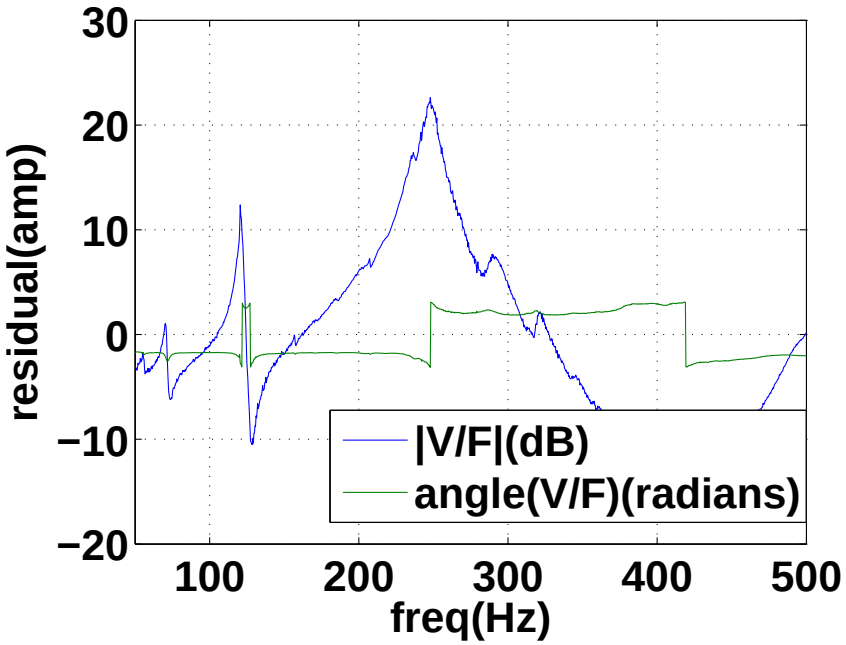

Figure 41. Measured body admittance from a gypsy guitar. The plot shows both the magnitude and phase at frequencies between 50Hz to 500Hz.

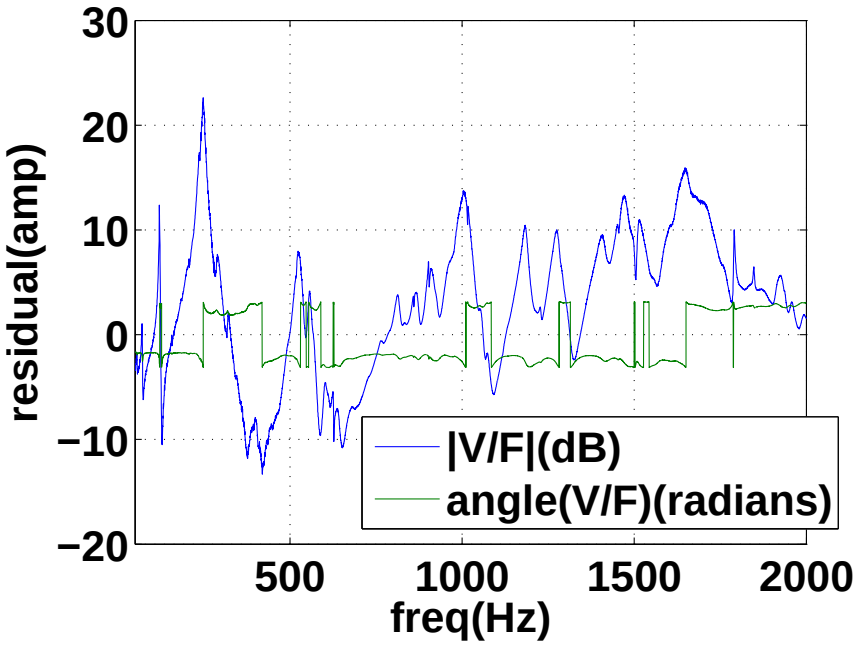

Figure 42. Measured body admittance from a gypsy guitar. The plot shows both the magnitude and phase at frequencies between 50Hz to 2000Hz.

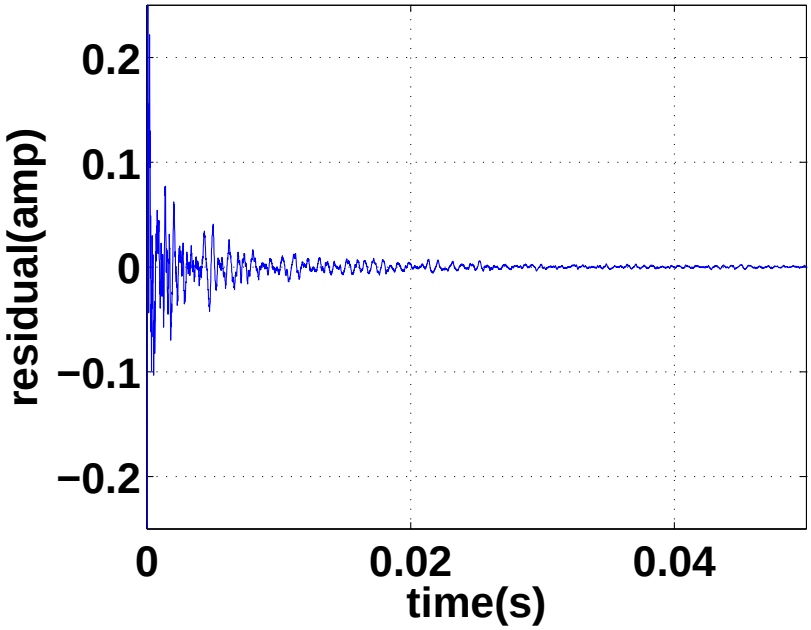

Figure 43. Time-domain signal of body response with peak at 120Hz removed.

4.3.1. Low-Order Filter Implementations

When latency and time-efficiency is a concern, the length of the body's impulse response becomes an issue. Therefore, in these systems, the impulse response of the body is approximated using lower-order filters and a modal synthesis model.

To effectively shorten the body's impulse response for systems requiring low-latency and low-memory needs, methods for removing the peaks are applied to leave a residual signal and a low-cost representation for the removed resonating peaks. This reduces the length of the original impulse response while representing the primary modes parametrically. The two general approaches for dealing with spectral peak removal include subtraction and inverse-filtering methods.

The two basic methods are as follows:

1. Complex Spectral Subtraction

$$H_r(z) = H(z) - \frac{b_0 + b_1 z^{-1}}{1 + a_1 z^{-1} + a_2 z^{-2}} \tag{15}$$

   where $H_r(z)$ corresponds to the shortened body impulse response while $H(z)$ corresponds to the measured body impulse response. The parameters to be estimated are the second-order filter coefficients b_0,b_1,a_1 and a_2.

   Complex Spectral Subtraction requires careful estimation of the phase, amplitude, frequency and bandwidth for peak removal. Furthermore, the resonators must run in parallel with the residual. Therefore, advantages obtained from Commuted Synthesis are lost as the approximated body impulse response model is not readily commutable with the string component of our physical model [45].

2. Inverse-Filtering

$$H_r(z) = H(z)(1 + a_1 z^{-1} + a_2 z^{-2}) \tag{16}$$

   where $H_r(z)$ again corresponds to the shortened body impulse response with $H(z)$ equal to the measured body impulse response. In this form, the residual signal is readily commutable with the string component of our physical model as resonators are factored instead of subtracted. Furthermore, estimating the coefficients of the filter for peak removal requires only the frequency and bandwidth of the peak and not the amplitude and phase as is required for Complex Spectral Subtraction [28].

Applying inverse-filtering as described above, we take the original body response shown in Figures 39 and 40, and remove the peak centered around 120Hz with a bandwidth of 10Hz. The residual, shown in Figure 43, is significantly shorter than the original response. Whereas in Figure 39, the response lasts for well over 50ms, in the residual signal, its amplitude oscillates near the noise floor at around 30ms. Figure 44 shows the spectrum, both magnitude and phase, of the response after inverse-filtering. Compared with Figure 41, the peak at 120Hz is completely removed.

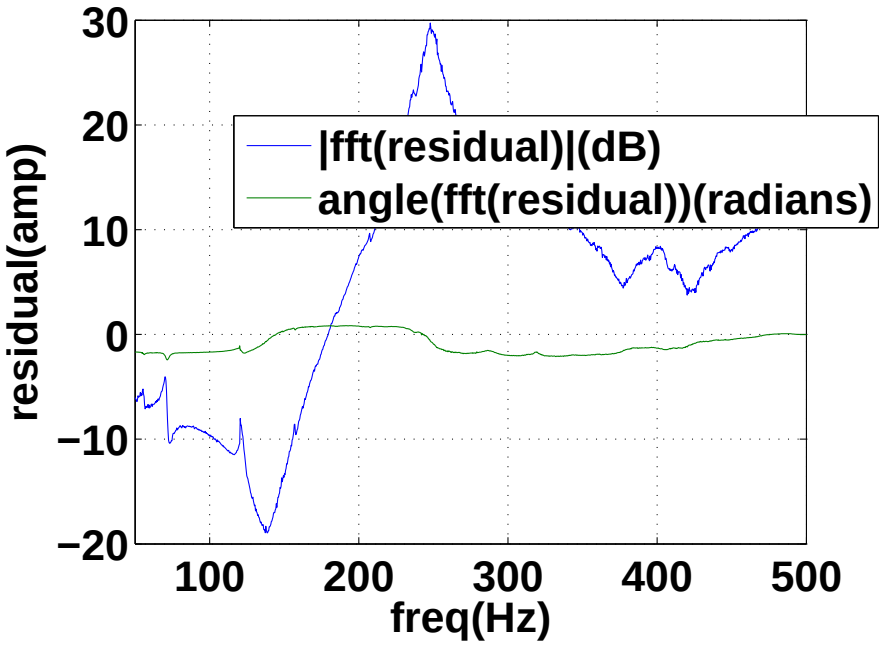

Figure 44. Spectral view, magnitude and phase, of body response with peak at 120Hz removed. Note that the response is now shorter compared with the original body response in Figure 39.

4.4. Radiated Sound Pressure

Using the measurements made in Section 3.4, we compute the point-to-point transfer function from the bridge to a single point in space. The resulting response is shown in Figure 45. The spectrum of the response is shown in Figure 46. Observing Figure 45, the length of the response is close to 2s long. Similar to the issues addressed with the body response, having to store the entire radiation response is costly. Therefore, the methods discussed for removing spectral peaks for the body response can be directly applied to the radiation response.

Similar to how the driving-point admittance is used to model the transfer of energy from the string to the bridge and vice-versa, the transfer function from a point-to-point measurement of acceleration of the top-plate to a point in space of pressure waves is used to affect the final sound output from our physical model.

4.4.1. Low-Order Filter Implementations

Using the same techniques described in Section 4.3.1, a reduction in the length of the resulting transfer function can be made by factoring out frequency components of the signal with the most energy and replacing them with low-order resonators. Both Complex Spectral Subtraction and inverse-filtering methods can be used. However, as done for the body response, we opt to use inverse-filtering. Figures 47 and 48 show the results of removing the peak centered around 127Hz with a bandwidth of 10Hz. Similar to the results obtained in Section 4.3.1, the radiation response is significantly shortened. As Figure 47 shows, the residual signal approaches 0 within 0.1s. Whereas, as shown in Figure 45 shows, the radiation response lasts for well over 1s. Furthermore, the peak in the spectral domain is entired removed. As shown in Figure 48, the peak is close to 60dBs lower than the peak at 127Hz in the original radiation response shown in Figure 46.

4.4.2. Interpolating Between Measurements

Because multiple plucks and excitations of a guitar can be made in one lab session, every note on the guitar can be recorded, thereby obtaining acceleration of the top-plate measurements and corresponding pressure waves for each note. For each note, computing the transfer function resulting from corresponding measurements yields correct radiation filter characteristics at the fundamental and harmonics of the note. Therefore, to obtain an averaged measurement radiation filter, using points from each notes' transfer functions corresponding to the fundamental frequency and harmonics, and interpolating between these known points also yields a psycho-acoustically accurate radiation filter.

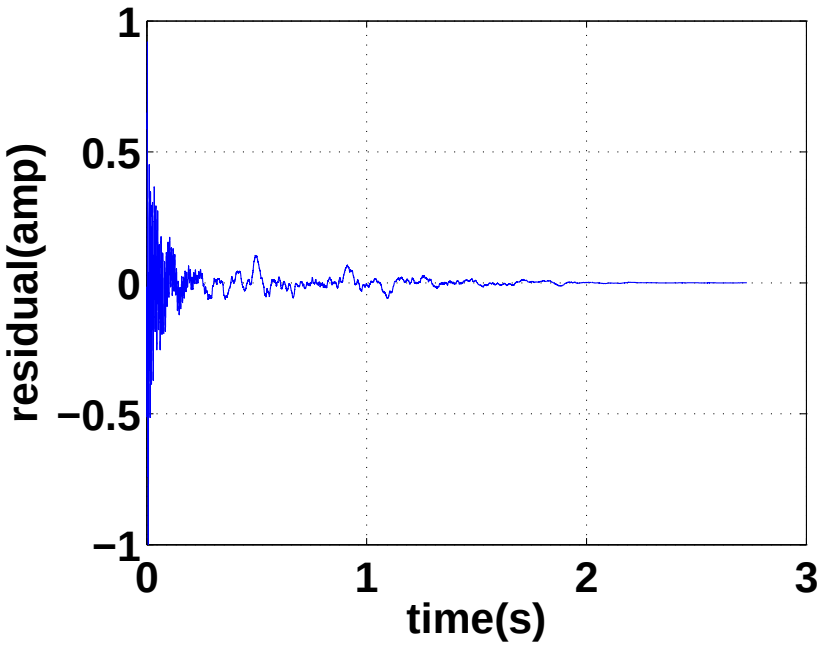

Figure 45. Measured radiation response from a gypsy guitar. The plot shows the signal from time 0 to 3s. Note the recordings were made in an anechoic chamber as described in Section 3.4.

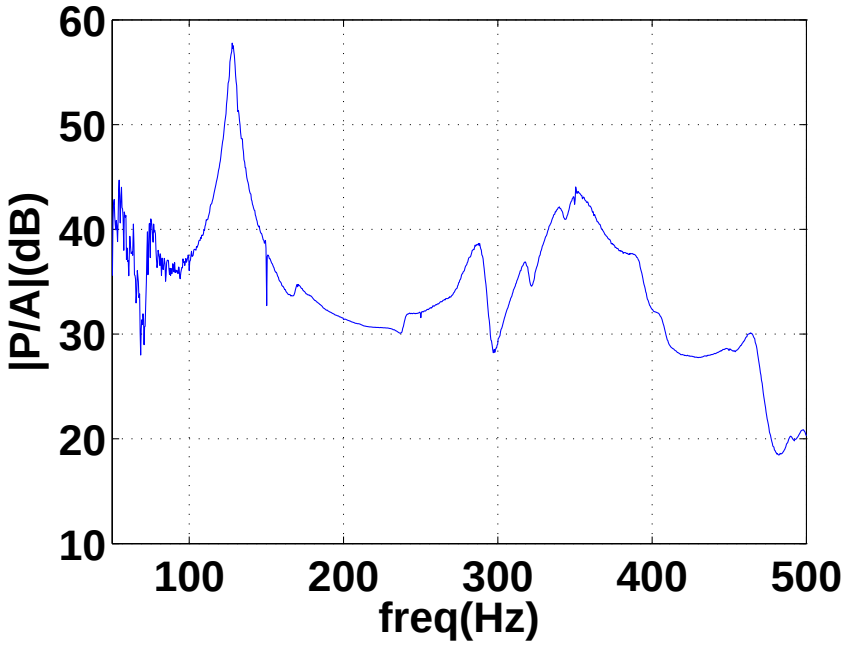

Figure 46. FFT of measured radiation response from a gypsy guitar. The plot shows the magnitude at frequencies between 50Hz to 500Hz.

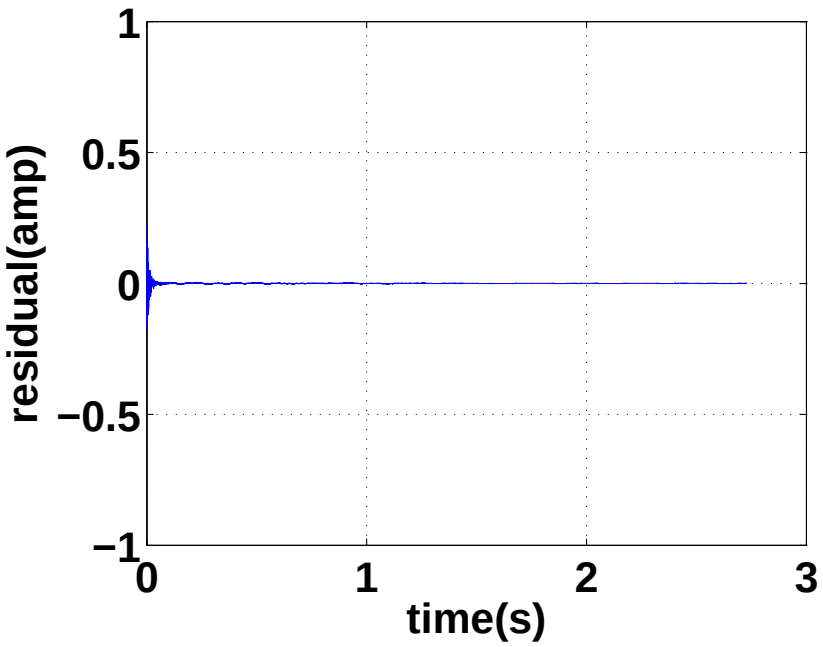

Figure 47. Removing the peak of the measured radiation response at 127Hz. The plot shows the signal from time 0 to 3s.

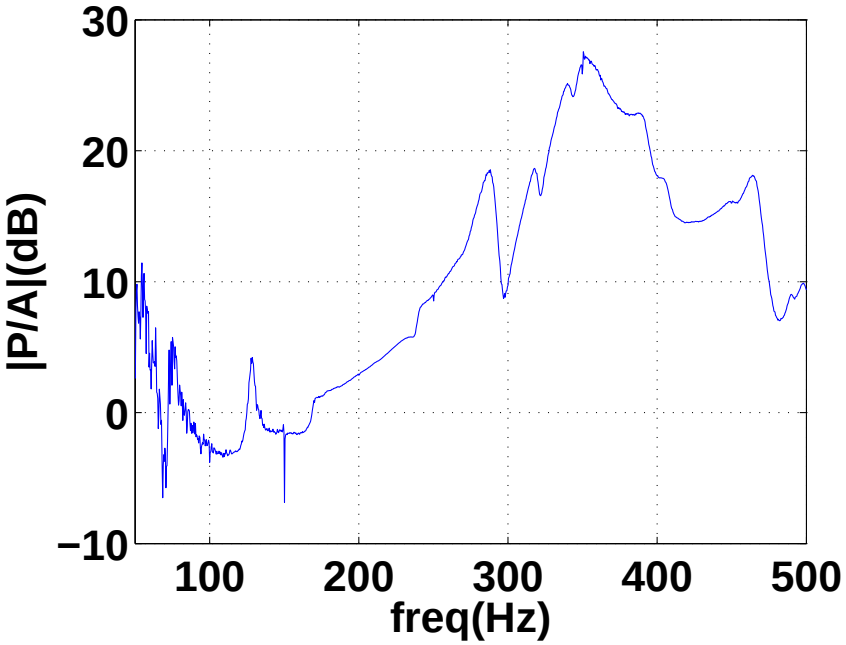

Figure 48. FFT of the measured radiation response with the peak at 127Hz removed. The plot shows the magnitude at frequencies between 50Hz to 500Hz.

5. REVIEW OF METHODS AND FINAL WORDS

In this reader, we discuss methods for decomposing a plucked stringed-instrument into modular components with individual properties that can be measured, modeled and calibrated. Within calibration, there are just a few overarching themes that should be mentioned to add cohesion to the methods presented. Inverse-filtering, used to reduce the order of measured body impulse responses and pressure radiation transfer functions, is extrememly useful in removing unwanted components within a given signal. However, there are conditions with which its method creates undesireable side-effects: nulls within the spectrum. As discussed in obtaining excitations for the digital waveguide string models in Section 4.1, using inverse-filtering causes nulls at the fundamental and its harmonics resulting in an anti-harmonic excitation. Since the excitation of the string is flat-spectrummed, inverse-filtering methods for excitation extraction does not present a psycho-acoustically pleasing solution.

6. REFERENCES

[1] J. l. R. d'Alembert, "Investigation of the curve formed by a vibrating string, 1747," *Acoustics: Historical and Philosophical Development* (R. B. Lindsay, ed.), pp. 119-123, Stroudsburg: Dowden, Hutchinson & Ross, 1973.

[2] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1975.

[3] J. O. Smith, "Introduction to Digital Filters with Audio Applications", *http://ccrma.stanford.edu/ jos/filters/, Sept. 2007*, online book.

[4] M.R. Schroeder, "Improved Quasi-Stereophony and Colorless Artificial Reverberation" *Journ. Audio Eng. Soc.*, 33:1061, 1961.

[5] J. A. Moorer, "About this reverberation business" *Computer Music Journal* vol. 3, no. 2, pp. 13-18, 1979.

[6] J. O. Smith, "Physical modeling using digital waveguides," *Computer Music Journal*, vol. 16, pp. 74-91, Winter 1992, special issue: Physical Modeling of Musical Instruments, Part I. http://ccrma.stanford.edu/ jos/pmudw/.

[7] A. Chaigne, "Viscoelastic properties of nylon guitar strings," *Catgut Acoustical Society Journal*, Series II, vol. 1, pp. 21-27, May 1991.

[8] J. O. Smith, "Efficient synthesis of stringed musical instruments," *Proceedings of the 1993 International Computer Music Conference*, Tokyo, pp. 64-71, Computer Music Association, 1993

[9] G. Weinreich, "Coupled piano strings," *Journal of the Acoustical Society of America*, vol. 62, pp. 1474-1484, Dec 1977

[10] N. Lee, A. Chaigne, J. O. Smith III, K. Arcas, "Measuring and Understanding the Gypsy Guitar," *International Symposium on Musical Acoustics*, Barcelona, Spain, September 9-12, 2007.

[11] J. Laroche, "Time and pitch scale modification of audio signals," *Applications of DSP to Audio & Acoustics* (M. Kahrs and K. Brandenburg, eds.), pp. 279-309, Boston/Dordrecht/London: Kluwer Academic Publishers, 1998

[12] J. O. Smith, "An allpass approach to digital phasing and flanging," *Proceedings of the 1984 International Computer Music Conference*, Paris, Computer Music Association, 1984, CCRMA Technical Report STAN-M-21, http://ccrma.stanford.edu/STANM/STANM/stanm14/.

[13] J. O. Smith, S. Serafin, J. Abel, and D. Berners, "Doppler simulation and the leslie," *Proceedings of the COST-G6 Conference on Digital Audio Effects* (DAFx-02), Hamburg, Germany, pp. 13-20, September 26 2002, http://www.dafx.de/.

[14] R. W. Schafer and L. R. Rabiner, "A digital signal processing approach to interpolation," *Proceedings of the IEEE*, vol. 61, pp. 692-702, June 1973.

[15] J. O. Smith and P. Gossett, "A flexible sampling-rate conversion method," *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, San Diego, vol. 2, (New York), pp. 19.4.1-19.4.2, IEEE Press, Mar. 1984, expanded tutorial and associated free software available at the Digital Audio Resampling Home Page: http://ccrma.stanford.edu/ jos/resample/.

[16] J. O. Smith and B. Friedlander, "Adaptive interpolated time-delay estimation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 21, pp. 180-199, Mar. 1985.

[17] T. I. Laakso, V. Välimäki, M. Karjalainen, and U. K. Laine, "Splitting the Unit Delay–Tools for Fractional Delay Filter Design," *IEEE Signal Processing Magazine*, vol. 13, pp. 30-60, Jan. 1996.

[18] W. Putnam and J. O. Smith, "Design of fractional delay filters using convex optimization," *Proceedings of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New Paltz, NY, (New York), IEEE Press, Oct. 1997, http://ccrma.stanford.edu/ jos/resample/optfir.pdf.

[19] D. A. Jaffe and J. O. Smith, "Extensions of the Karplus-Strong plucked string algorithm," *Computer Music Journal*, vol. 7, no. 2, pp. 56-69, 1983,

[20] K. Karplus and A. Strong, "Digital synthesis of plucked string and drum timbres," *Computer Music Journal*, vol. 7, no. 2, pp. 43-55, 1983,

[21] A. Chaigne, "On the use of finite differences for musical synthesis. application to plucked stringed instruments," *Journal d'Acoustique*, vol. 5, no. 2, pp. 181-211, 1992.

[22] J. O. Smith and S. A. Van Duyne, "Commuted piano synthesis," *Proceedings of the 1995 International Computer Music Conference*, Banff, pp. 319-326, Computer Music Association, 1995, http://ccrma.stanford.edu/ jos/pdf/svd95.pdf

[23] N. H. Fletcher and T. D. Rossing, *The Physics of Musical Instruments*, 2nd Edition, New York: Springer Verlag, 1998.

[24] M. E. Van Valkenburg, *Introduction to Modern Network Synthesis*, New York: John Wiley and Sons, Inc., 1960.

[25] R. W. Newcomb, *Linear Multiport Synthesis*, New York: McGraw-Hill, 1966.

[26] V. Belevitch, *Classical Network Theory*, San Francisco: Holden Day, 1968.

[27] M. Karjalainen, V. Välimäki, and Z. Jnosy, "Towards high-quality sound synthesis of the guitar and string instruments," *Proceedings of the 1993 International Computer Music Conference*, Tokyo, pp. 56-63, Computer Music Association, Sept. 10-15 1993, http: //www.acoustics.hut.fi/ vpv/publications/icmc93-guitar.htm.

[28] J.O. Smith, *Physical Audio Signal Processing*, August 2007 Edition, http://ccrma.stanford.edu/ jos/pasp/, online book, accessed 11/21/07.

[29] L. V. King On the Acoustic Radiation Pressure on Spheres *Proceedings of the Royal Society of London.* Series A, Mathematical and Physical Sciences, Vol. 147, No. 861. (Nov. 15, 1934), pp. 212-240.

[30] O. Christensen and R.B. Vistisen, A simple model for low frequency guitar function, *J. Acoust. Soc. Am.*, vol. 68, pp. 758-766, 1980.

[31] B.E. Richardson, T. J. W. Hill and S.J. Richardson, Input admittance and sound field measurements of ten classical guitars, *Proceedings of the Institute of Acoustics* 2002, vol. 24(2), pp. 1-10, IOA.

[32] J. P. Dalmont and A. M. Bruneau, Acoustic impedance measurements: plane-wave mode and first helical mode contributions. *Journal of the Acoustical Society of America* 91 (1992), pp. 3026-3033.

[33] N. Lee, Z. Duan, J. O. Smith III, Excitation Extraction for Guitar Tones, *Proceedings of the International Computer Music Conference* (ICMC07), Copenhagen, Denmark, August 27-31, 2007.

[34] Y. Hua and T. Sarkar, "Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise", *IEEE Trans. Acoustics, Speech, Signal Processing*, vol. ASSP-38, no. 5, pp. 814–824, May 1990.

[35] J. Laroche and J. L. Meillier, Multichannel excitation filter modeling of percussive sounds with application to the piano, *IEEE Trans. Speech and Audio Processing*, vol. 2, no. 2, pp. 329-344, Apr. 1994.

[36] T. Tolonen, Model-based analysis and resynthesis of acoustic guitar tones, *Masters thesis*, Helsinki University of Technology, Espoo, Finland, Jan. 1998, Report 46, Laboratory of Acoustics and Audio Signal Processing.

[37] V. Välimäki and T. Tolonen, "Development and calibration of a guitar synthesizer", *J. Audio Eng. Soc.*, vol. 46, no. 9, Sept. 1998, Proceedings of the 103rd Convention of the Audio Engineering Society, New York, Sept. 1997.

[38] X. Serra, "Musical sound modeling with sinusoids plus noise", in *Musical Signal Processing*, C. Roads, S. Pope, A. Picialli, and G. De Poli, Eds., pp. 91–122. Swets & Zeitlinger Publishers, Netherlands, June 1997.

[39] N. Laurenti and G. D. Poli, A method for spectrum separation and envelope estimation of the residual in spectrum modeling of musical sound, *Proc. Int. Conf. Digital Audio Effects* (DAFx-00), Verona, Italy, 2000, http://www.dafx.de/.

[40] J. O. Smith and X. Serra, "PARSHL: A program for the analysis/synthesis of inharmonic sounds based on a sinusoidal representation", in *Proc. 1987 Int. Computer Music Conf., Champaign-Urbana*. 1987, Computer Music Association. Expanded version available on-line at `http://ccrma.stanford.edu/~jos/parshl/`.

[41] N. Lee, R. Cassidy, J. O. Smith III, Use of Energy Decay Relief (EDR) to Estimate Partial-Overtone Decay-Times in a Freely Vibrating String, Invited Paper at *The Musical Acoustics Sessions at the Joint ASA-ASJ Meeting*, Honolulu, Hawaii (2006)

[42] J. O. Smith, Techniques for Digital Filter Design and System Identification with Application to the Violin, PhD thesis, Elec. Engineering Dept., Stanford University (CCRMA), June 1983, CCRMA Technical Report STAN-M-14, http://ccrma.stanford.edu/STANM/STANM/stanm14/.

[43] B. Bank and V. Välimäki, "Robust loss filter design for digital waveguide synthesis of string tones," *IEEE Signal Processing Letters*, vol. 10, pp. 18-20, January 2003.

[44] C. Lambourg and A. Chaigne, "Measurements and modeling of the admittance matrix at the bridge in guitars," *Proceedings of the Stockholm Musical Acoustics Conference* (SMAC-93), pp. 448-453, July 1993.

[45] K. Bradley, M. H. Cheng, and V. L. Stonick, "Automated analysis and computationally efficient synthesis of acoustic guitar strings and body," *Proceedings of the IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, New Paltz, NY, (New York), pp. 238-241, IEEE Press, Oct. 1995, Session 9a, paper 7, 4 pages.