

Efficient Synthesis of Stringed Musical Instruments

Julius O. Smith III (jos@ccrma.stanford.edu)
Center for Computer Research in Music and Acoustics (CCRMA)
Department of Music, Stanford University
Stanford, California 94305

Abstract

Techniques are described for reducing complexity in stringed instrument simulation for purposes of digital synthesis. These include commuting losses and dispersion to consolidate them into a single filter, replacing body resonators by look-up tables, simplified bow-string interaction, and single-filter, multiply-free coupled strings implementation.

Contents

1	Digital Waveguide Theory	2
2	The Terminated String	4
3	Simplified Body Filters	5
4	Simplified Bowed Strings	8
5	Coupled Strings	10
6	Summary	14
7	Appendix	14

1 Digital Waveguide Theory

This section summarizes the *digital waveguide* model for vibrating strings. Further details can be found in [Smith 1992].

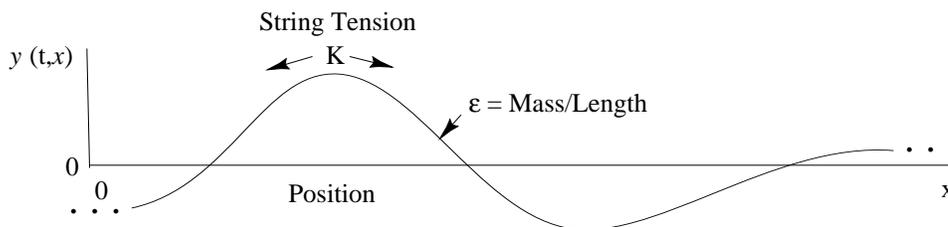


Figure 1: The ideal vibrating string.

The *wave equation* for the ideal (lossless, linear, flexible) vibrating string, depicted in Fig. 1, is given by

$$Ky'' = \epsilon \ddot{y}$$

where

$$\begin{array}{ll} K \triangleq & \text{string tension} & y \triangleq & y(t, x) \\ \epsilon \triangleq & \text{linear mass density} & \dot{y} \triangleq & \frac{\partial}{\partial t} y(t, x) \\ y \triangleq & \text{string displacement} & y' \triangleq & \frac{\partial}{\partial x} y(t, x) \end{array}$$

The same wave equation applies to any displacement along one dimension in any perfectly elastic medium. We refer to the general class of such media as *one-dimensional waveguides*. Extension to two and more dimensions is described elsewhere in this proceedings [Van Duyne and Smith 1992].

It can be readily checked that the wave equation is solved by any string shape which travels to the left or right with speed $c = \sqrt{K/\epsilon}$. (But note that the derivation of the wave equation assumes the string slope is much less than 1 at all times and positions.) If we denote right-going traveling waves by $y_r(x - ct)$ and left-going traveling waves by $y_l(x + ct)$, where y_r and y_l are arbitrary twice-differentiable functions, then the general class of solutions to the lossless, one-dimensional, second-order wave equation can be expressed as

$$y(x, t) = y_r(x - ct) + y_l(x + ct)$$

Sampling the traveling-waves gives

$$\begin{aligned} y(t_n, x_m) &= y_r(t_n - x_m/c) + y_l(t_n + x_m/c) \\ &= y_r(nT - mX/c) + y_l(nT + mX/c) \\ &= y_r[(n - m)T] + y_l[(n + m)T] \end{aligned}$$

Since T multiplies all arguments, we suppress it by defining

$$y^+(n) \triangleq y_r(nT) \quad y^-(n) \triangleq y_l(nT)$$

The “+” superscript denotes a traveling-wave component propagating to the right, and “-” denotes propagation to the left. Finally, the left- and right-going traveling waves must be summed to produce a physical output according to the formula

$$y(t_n, x_m) = y^+(n - m) + y^-(n + m)$$

The appendix shows a linear wave equation with constant coefficients, of any order, admits a decaying, dispersive, traveling-wave solution. Even-order time derivatives give rise to dispersion and odd-order time derivatives correspond to losses; higher order spatial derivatives give rise to multiple solutions of the same kind. The corresponding digital simulation of an arbitrarily long (undriven and unobserved) section of medium can be simplified via commutativity to at most two pure delays and at most two linear, time-invariant filters. In dimensions higher than one, these remarks apply to any given direction of traveling-wave propagation.

Since every linear, time-invariant filter can be expressed as a zero-phase filter in cascade with an allpass filter, we may factor the filter into its lossy part and its dispersive part. The zero-phase factor implements frequency-dependent gain (damping in a digital waveguide), and the allpass part gives frequency-dependent delay, (dispersion in a digital waveguide). A digital simulation diagram appears in Fig. 2.

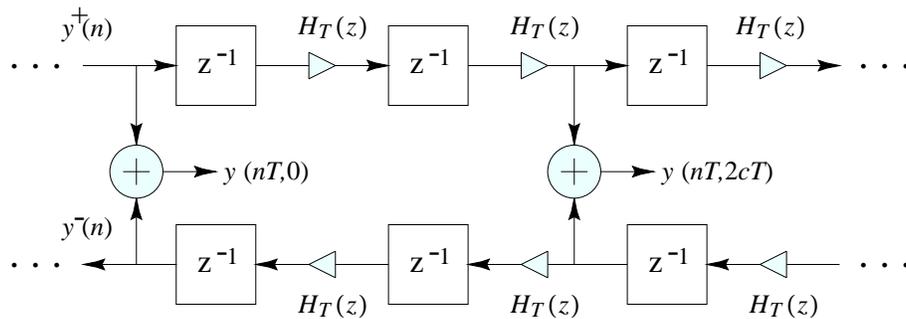


Figure 2: Discrete simulation of the ideal, linear, lossy, dispersive, digital waveguide.

The simulation of the traveling-waves is exact, in principle, at the sampling positions and instants, even though losses and dispersion are admitted in the wave equation. Note also that the losses which are *distributed* in the continuous solution have been consolidated, or *lumped*, at discrete intervals of cT meters in the simulation. The filter $H_T(z)$ *summarizes* the distributed filtering incurred in one sampling interval. The lumping of distributed filtering does not introduce an approximation error at the sampling points. Furthermore, bandlimited interpolation can

yield arbitrarily accurate reconstruction between samples [Smith and Gossett 1984]. The main restriction is that all initial conditions and excitations be bandlimited to half the sampling rate.

It is usually possible to realize vast computational savings in waveguide simulation by *commuting losses out of unobserved and undriven sections of the medium and consolidating them at a minimum number of points*. Because the digital simulation is linear and time invariant (given constant medium parameters), and because linear, time-invariant elements commute, the diagram in Fig. 3 is exactly equivalent (to within numerical precision) to the previous diagram in Fig. 2. Each per-sample filter $H_T(z)$ is commuted with delay elements and combined with other filters until an input or output is encountered which inhibits further migration. Filters can also be pushed through nodes in the diagram to achieve further simplifications in some cases.

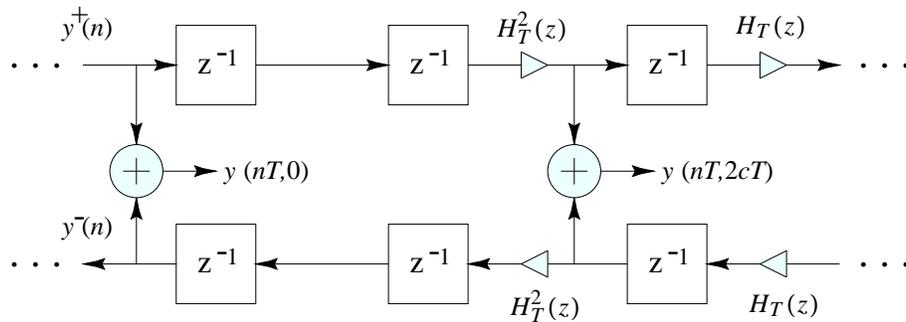


Figure 3: General linear digital waveguide with commuted loss/dispersion filters.

2 The Terminated String

Using the above simplification principles, it is possible to commute the elements of a rigidly terminated, dispersive, lossy string into the form shown in Fig. 4, provided that the string is to be excited by initial conditions and the output signal is taken to be a traveling-wave component. In this case, the losses and dispersion are lumped at a single point in the round-trip travel along the string. When the loop filter is a two-point average $(1+z^{-1})/2$, and when the initial conditions used to “pluck” the string are taken to be random numbers, the well known Karplus-Strong algorithm for string and drum sounds is obtained [Karplus and Strong 1983, Jaffe and Smith 1983].

In a general physical model, the loop filter is determined by the cascade of (1) the filtering experienced by a traveling wave in traversing the string twice, and (2) the reflection transfer functions of the two terminations.

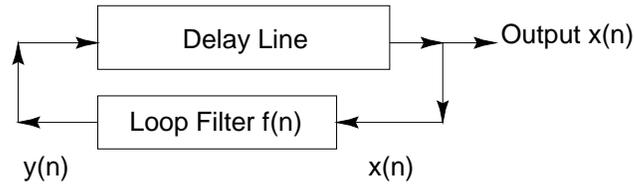


Figure 4: The rigidly terminated, linear string.

If the wave impedance of the string is $R = \sqrt{K\epsilon}$, and the bridge driving-point impedance is $R_b(z)$, then the reflection transfer function at the bridge is given by

$$S_b(z) = \frac{F^-(z)}{F^+(z)} = \frac{R_b(z) - R}{R_b(z) + R}$$

for force waves, and $-S_b(z)$ for velocity waves. Because the bridge is passive, $R_b(z)$ is *positive real*, [Van Valkenburg 1960], i.e.,

- (1) $R_b(z)$ is real when z is real.
- (2) $|z| \geq 1 \Rightarrow \text{re} \{R_b(z)\} \geq 0$.

This implies $S_b(z)$ is a *Schur function*, i.e., $S_b(z) \leq 1$ for $|z| \geq 1$. Reflection filters associated with passive, finite-order impedances always have an equal number of poles and zeros, as can be seen from the above expression. If the bridge termination is lossless, its impedance $R_b(z)$ is purely reactive and the reflection filter $S_b(z)$ becomes allpass. Typically, the reflection filter has gain less than but close to 1 at all frequencies, and the gain is smallest at frequencies where there is strong coupling with a bridge or body resonance.

3 Simplified Body Filters

In a complete stringed musical instrument, such as a guitar, the string couples via the bridge into a resonating “body” which is needed for coupling to the surrounding air, and which imposes a frequency response of its own on the radiated sound. In addition, spectral characteristics of the string excitation affect the radiated sound. Thus, we have the components shown in Fig. 5.



Figure 5: Schematic diagram of a stringed musical instrument.

Because the string and body are approximately linear and time-invariant, we may *commute the string and resonator*, as shown in Fig. 6.



Figure 6: Equivalent diagram in the linear, time-invariant case.

The excitation can now be *convolved* with the resonator impulse response to provide a single, aggregate, excitation table, as depicted in Fig. 7.

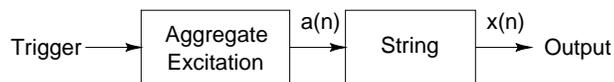


Figure 7: Use of an aggregate excitation given by the convolution of original excitation with the resonator impulse response.

In the simplest case, the string is “plucked” using the (half-windowed) impulse response of the body.

An example of an excitation is the force applied by a pick or a finger at some point, or set of points, along the string. The input force per sample at each point divided by $4R$ gives the velocity to inject additively at that point in both traveling-wave directions. (The factor of 4 comes from splitting the injected velocity into two traveling-wave components, and from the fact that two string end-points are being driven.) Equal injection in the left- and right-going directions corresponds to an excitation force which is stationary with respect to the string.

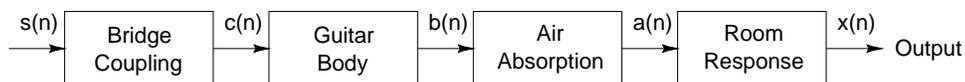


Figure 8: Possible components of a guitar resonator.

In a practical instrument, the “resonator” is determined by the choice of output signal in the physical scenario, and it generally includes filtering downstream of the body itself, as shown in Fig. 8. A typical example for the guitar or violin would be to choose the output signal at a point a few feet away from the top plate of the body. In practice, such a signal can be measured using a microphone held at the desired output point and recording the response at that point to the striking of the bridge with a force hammer. It is useful to record simultaneously the output of an accelerometer mounted on the bridge in order to also obtain experimentally the driving-point impedance at the bridge. In general, it is desirable to choose the output close to the instrument so as to keep the resonator response as short as possible. The resonator components need to be linear and time invariant, so they will be commutative with the string and combinable with the string excitation signal via convolution.

The string should also be linear and time invariant in order to be able to commute it with the generalized resonator. However, the string is actually the least linear element of most stringed musical instruments, with the main effect of nonlinearity being a slight increase of the fundamental vibration frequency with amplitude. A secondary effect is to introduce coupling between the two polarizations of vibration along the length of the string. In practice, however, the string can be considered sufficiently close to linear to permit commuting with the body. The string is also time varying in the presence of vibrato, but this too can be neglected in practice. While commuting a live string and resonator may not be identical mathematically, the sound is substantially the same.

There are various options when combining the excitation and resonator into an aggregate excitation, as shown in Fig. 7. For example, a wave-table can be prepared which contains the convolution of a particular point excitation with a particular choice of resonator. Perhaps the simplest choice of excitation is the impulse signal. Physically, this would be natural when the wave variables in the string are taken to be acceleration waves for a plucked string; in this case, an ideal pluck gives rise to an impulse of acceleration input to the left and right in the string at the pluck point. If loss of perceived pick position is unimportant, the impulse injection need only be in a single direction. (The comb filtering which gives rise to the pick-position illusion can be restored by injecting a second, negated impulse at a delay equal to the travel time to and from the bridge.) In this simple case of a single impulse to pluck the string, the aggregate excitation is simply the impulse *response* of the resonator. Many excitation and resonator variations can be simulated using a collection of aggregate excitation tables. It is useful to provide for interpolation of excitation tables so as to provide intermediate points along a parameter dimension. In fact, all the issues normally associated with sampling synthesis arise in the context of the string excitation table. A disadvantage of combining excitation and resonator is the loss of multiple output signals from the body simulation, but the timbral effects arising from the mixing together of multiple body outputs can be obtained via a mixing of corresponding excitation tables.

If the aggregate excitation is too long, it may be shortened by a variety of techniques. It is good to first convert the final excitation $a(n)$ in Fig. 7 to *minimum phase* so as to provide the maximum shortening consistent with the original magnitude spectrum. Secondly, $a(n)$ can be “half-windowed” using the right wing any window function typically used in spectrum analysis. An interesting choice is the exponential window, since it has the interpretation of increasing the resonator damping in a uniform manner, i.e., all the poles and zeros of the resonator are contracted radially in the z plane by the same factor.

4 Simplified Bowed Strings

The method of the previous section can be extended to bowed strings in an efficient way. The “leaning sawtooth” waveforms observed by Helmholtz for steady state bowed strings can be obtained by periodically “plucking” the string in only one direction along the string. In principle, a traveling impulsive excitation is introduced into the string in the right-going direction each period for a “down bow” and in the left-going direction for an “up bow.” This simplified bowing simulation works best for smooth bowing styles in which the notes have slow attacks. More varied types of attack can be achieved using the more physically accurate McIntyre-Woodhouse theory [Smith 1987].

Commuting the string and resonator means that the string is now plucked by a *periodically repeated resonator impulse response*. A nice simplified vibrato implementation is available by varying the impulse-response retriggering period, i.e., the vibrato is implemented in the excitation oscillator and not in the delay loop. The string loop delay need not be modulated at all. While this departs from being a physical model, the vibrato quality is satisfying and qualitatively similar to that obtained by a rigorous physical model. Figure 9 illustrates the overall block diagram of the simplified bowed string and its commuted and response-excited versions.

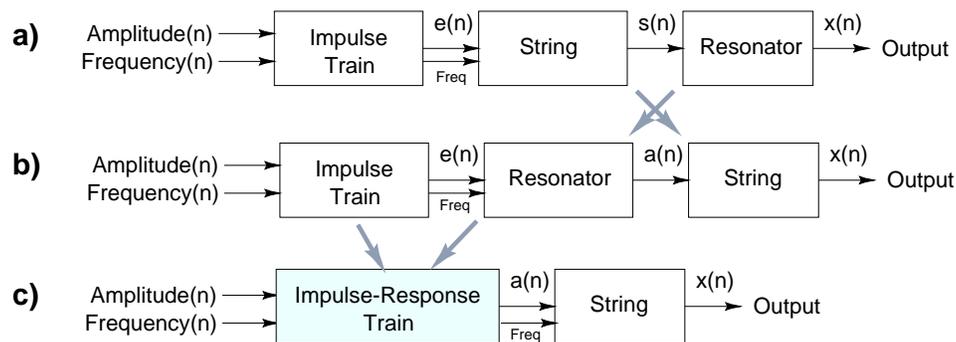


Figure 9: a) The simplified bowed string, including amplitude, pitch, and vibrato controls. The frequency control is also used by the string. b) Equivalent diagram with resonator and string commuted. c) Equivalent diagram in which the resonator impulse response is played into the string each pitch period.

In current technology, it is reasonable to store one recording of the resonator impulse response in digital memory as one of many possible *string excitation tables*. The excitation can contribute to many aspects of the tone to be synthesized, such as whether it is a violin or a cello, the force of the bow, and where the bow is playing on the string. Also, graphical equalization and other time-invariant filtering can be provided in the form of alternate excitation-table choices.

During the synthesis of a single bowed-string tone, the excitation signal is played into the string quasi-periodically. Since the excitation signal is typically longer than one period of the tone, it is necessary to either (1) interrupt the excitation playback to replay it from the beginning, or

(2) start a new playback which overlaps with the playback in progress. Variant (2) requires a separate incrementing pointer and addition for each instance of the excitation playback; thus it is more expensive, but it is preferred from a quality standpoint.

Of course, ordinary wavetable synthesis or any other type of synthesis can also be used as an excitation signal in which case the string loop behaves as a pitch-synchronous comb filter following the wavetable oscillator. Interesting effects can be obtained by slightly detuning the wavetable oscillator and delay loop; tuning the wavetable oscillator to a harmonic of the delay loop can also produce an ethereal effect.

The externally excited, filtered delay loop can be used also to simulate wind and other musical instruments. In fact, *any quasi-periodic tone* can be approximated using an appropriate excitation signal (which may be varied over time) together with some loop filter (which also may be varied over time). The fact that the delay line is approximately one period in length restricts application of this type of structure to quasi-periodic tones. However, aperiodic tones which can be well approximated by a superposition of a few quasi-periodic tones can be synthesized using multiple delay loops added together in parallel and excited by common or separate excitations. Thus, piano, marimba, and glockenspiel can be approximated, for example. For wind instruments, a filtered, enveloped noise excitation is needed. In summary, the externally excited, filtered delay loop can be viewed as an efficient compression technique for arbitrary quasi-periodic signals with musically desirable parameters.

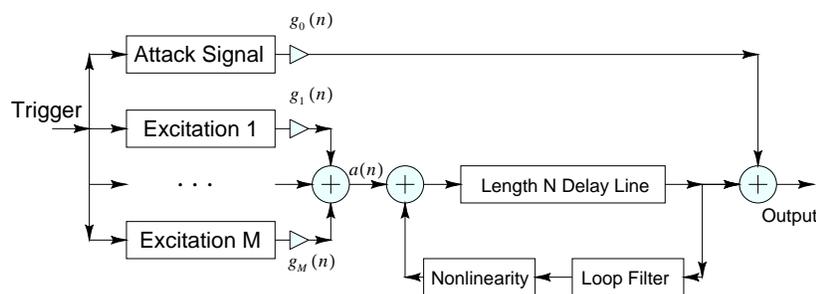


Figure 10: Generalized filtered delay loop synthesis.

Figure 10 illustrates a more general version of the table-excited, filtered delay loop synthesis system. The generalizations help to obtain a wider class of timbres. The multiple excitations summed together through time-varying gains provide for timbral evolution of the tone. For example, a violin can transform smoothly into a cello, or the bow can move smoothly toward the bridge by interpolating among two or more tables. Alternatively, the tables may contain “principal components” which can be scaled and added together to approximate a wider variety of excitation timbres. An excellent review of multiple wavetable synthesis appears in [Horner *et al.* 1993]. The nonlinearity is useful for obtaining distortion guitar sounds and other interesting evolving timbres.

Finally, the “attack signal” path around the string has been found to be useful for reducing the cost of implementation: the highest frequency components of a struck string, say, tend to emanate immediately from the string to the resonator with very little reflection back into the

string (or pipe, in the case of wind instrument simulation). Injecting them into the delay loop increases the burden on the loop filter to quickly filter them out. Bypassing the delay loop altogether alleviates requirements on the loop filter and even allows the filtered delay loop to operate at a lower sampling rate; in this case, a signal interpolator would appear between the string output and the summer which adds in the scaled attack signal in Fig. 10. For example, it was found that the low E of an electric guitar (Gibson Les Paul) can be synthesized quite well using a filtered delay loop running at a sampling rate of 3 kHz. (The pickups do not pick up much energy above 1.5 kHz.) Similar savings can be obtained for any instrument having a high-frequency content which decays much more quickly than its low-frequency content.

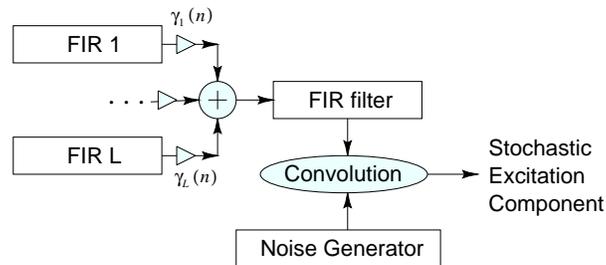


Figure 11: Example of a filtered noise excitation implementation.

For good generality, at least one of the excitation signals should be a filtered noise signal. An example implementation is shown in Fig. 11. In this example, there is a free running bandlimited noise generator which is filtered by a finite impulse response (FIR) digital filter. The filter coefficients are computed in real time as a linear combination of a set of fixed FIR coefficient sets stored in ROM. A recursive filter may also be used, in which case ladder/lattice forms can be used so that the coefficients can be interpolated without stability problems. In a simple implementation, only two gains might be used, allowing simple interpolation from one filter to the next, and providing an overall amplitude control for the noise component of the excitation signal.

5 Coupled Strings

In stringed musical instruments, coupling phenomena cannot be ignored. Coupling effects include amplitude modulation of partial amplitude envelopes due to “beating” between two or more coupled modes, two-stage decay (a fast decay followed by a slower decay), or “aftersound” [Weinreich 1977]. Physically, significant coupled-string phenomena result from inter-string coupling, coupling between the horizontal and vertical polarizations of vibration on one string, and between the string and body resonances.

The simplest simulation of coupled strings is obtained by simply summing two or more slightly detuned strings. More realistic string coupling involves actual signal flow from each coupled string to all others.

“Efficient Synthesis of Stringed Musical Instruments,” J.O. Smith, Proc. Int. Computer Music Conf. (ICMC-93), pp. 64–71, Tokyo, 1993, Computer Music Association.

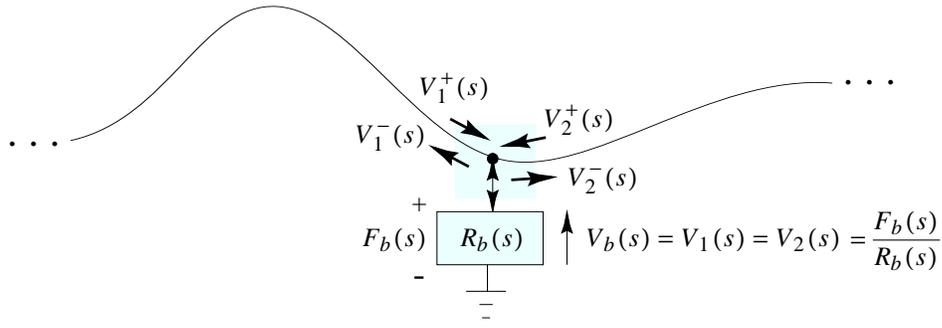


Figure 12: Two strings terminated at a common bridge impedance.

A diagram for the two-string case is shown in Fig. 12. This situation is a special case of the *loaded waveguide junction* [Smith 1987], with the number of waveguides being $N = 2$, and the junction load being the transverse driving-point impedance $R_b(s)$ where the string drives the bridge. For a direct derivation, we need only observe that (1) the string velocities of each string endpoint must each be equal to the velocity of the bridge, or $v_1 = v_2 = v_b$, and (2) the sum of forces of both strings equals the force applied to the bridge: $f_b = f_1 + f_2$. The bridge impedance relates the force and velocity of the bridge via $F_b(s) = R_b(s)V_b(s)$. Expanding into traveling wave components in the Laplace domain, we have

$$\begin{aligned}
 R_b(s)V_b(s) &= F_b(s) = F_1(s) + F_2(s) \\
 &= [F_1^+(s) + F_1^-(s)] + [F_2^+(s) + F_2^-(s)] \\
 &= R_1\{V_1^+(s) - [V_b(s) - V_1^+(s)]\} \\
 &\quad + R_2\{V_2^+(s) - [V_b(s) - V_2^+(s)]\}
 \end{aligned}$$

or

$$V_b(s) = H_b(s)[R_1V_1^+(s) + R_2V_2^+(s)]$$

where R_i is the wave impedance of string i , and

$$H_b(s) \triangleq \frac{2}{R_b(s) + R_1 + R_2}$$

Thus, in the time domain, the incoming velocity waves are scaled by their respective wave impedances, summed together, and filtered according to the transfer function $H_b(s) = 2/[R_b(s) + R_1 + R_2]$ to obtain the velocity of the bridge $v_b(t)$.

Given the filter output $v_b(t)$, the outgoing traveling velocity waves are given by

$$\begin{aligned}
 v_1^-(t) &= v_b(t) - v_1^+(t) \\
 v_2^-(t) &= v_b(t) - v_2^+(t)
 \end{aligned}$$

Thus, the incoming waves are subtracted from the bridge velocity to get the outgoing waves.

Since $V_2^-(s) = H_b(s)R_1V_1^+(s) = H_b(s)F_1^+(s)$ when $V_2^+(s) = 0$, and vice versa exchanging strings 1 and 2, H_b may be interpreted as the *transmission admittance filter* associated with the bridge coupling. It can also be interpreted as the bridge admittance transfer function from every string, since its output is the bridge velocity resulting from the sum of incident traveling force waves.

A general coupling matrix contains a filter transfer function in each entry of the matrix. For N strings, each conveying a single type of wave (e.g., horizontally polarized), the general linear coupling matrix would have N^2 transfer-function entries. In the present formulation, only one transmission filter is needed, and it is shared by all the strings meeting at the bridge.

The above sequence of operations is formally similar to the *one multiply scattering junction* frequently used in digital lattice filters [Markel and Gray 1976]. In this context, it would be better termed the “one-filter scattering termination.”

When the two strings are identical (as would be appropriate in a model for coupled piano strings), the computation of bridge velocity simplifies to

$$V_b(s) = H_b(s)[V_1^+(s) + V_2^+(s)]$$

where $H_b(s) \triangleq 2/[2 + R_b(s)/R]$ is the *velocity* transmission filter. In this case, the incoming velocities are simply summed and fed to the transmission filter which produces the bridge velocity at its output. A commuted simulation diagram appears in Fig. 13.

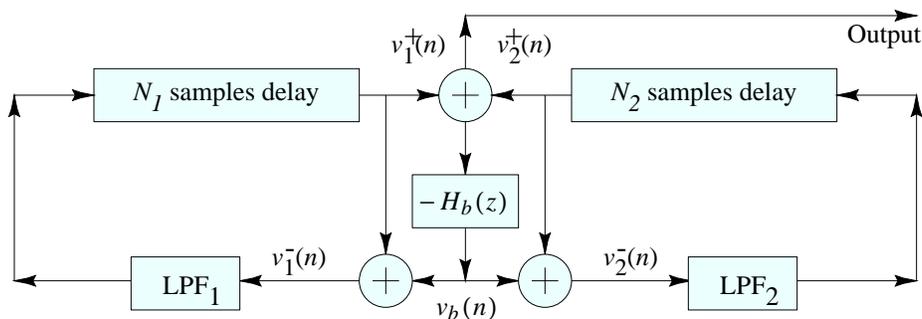


Figure 13: General linear coupling of two equal-impedance strings using a common bridge filter.

Since $R_b(z)$ is positive real, it is readily verified that

$$\left| 2H_b(e^{j\omega T}) - 1 \right| \leq 1$$

which restricts the set of coupling filters to those having frequency response values in the circle of radius $1/2$ centered at $z = 1/2$ in the complex plane. If the two coupled strings are taken to be lossless (e.g., two pure delay loops), then this constraint becomes the stability condition for

the overall system. If the amplitude and phase response of the filter are denoted $G(\omega)$ and $\theta(\omega)$, respectively, the passivity constraint may be written in the form

$$\cos[\theta(\omega)] \geq G(\omega)$$

Thus, the gain may approach unity only at frequencies where the phase approaches zero. In no case may the absolute value of the phase exceed 90 degrees, nor may the gain exceed 1 at any frequency. If the phase does approach plus or minus 90 degrees, the gain must approach zero also. The real part of the frequency response is always positive, and it may approach zero only if the imaginary part (hence gain) also approaches zero.

If the transmission filter H_b is taken to be a real, frequency-independent gain G , corresponding to a “resistive bridge termination,” the passivity constraint becomes simply

$$0 \leq G \leq 1$$

Such a set of resistive bridge couplings may be realized *without multiplies* by using gain values of the form

$$G = 2^{-K}, \quad K = 0, 1, 2, \dots$$

The case $G = 1$ corresponds to a zero bridge impedance which means the two strings simply fuse into one long ideal string. The case $G = 0$ corresponds to an infinitely rigid bridge, in which case the two strings are isolated from one another. Since realistic bridges are close to rigid, we desire many settings in the vicinity of $G = 0$, and the “right-shift” $G = 2^{-K}$ has this property.

Another passive, multiply-free, transmission filter is any filter having a transfer function of the form

$$H_b(z) = 2^{-K}(1 + z^{-1}), \quad K = 1, 2, \dots$$

Thus, the right-shifter is augmented by a unit-sample delay and a summer. In this case, the bridge appears more rigid at high frequencies, behaving like a mass. Spring-like bridges can be implemented using a transmission filter of the form $H_b(z) = 2^{-K}(1 - z^{-1})$, $K = 1, 2, \dots$. These are one-zero filters. Corresponding multiply-free one-pole versions are $H_b(z) = 2^{-K}/(1 - z^{-1})$ for a mass-like bridge and $H_b(z) = 2^{-K}/(1 + z^{-1})$ for a spring-like bridge.

Any passive transmission filter can be cascaded with any resistive loss. Also, one mass-like and one spring-like transmission filter as defined above can be cascaded. However, instability can result if two mass or two spring filters are used in cascade. For higher orders, it is necessary to go to second-order sections whose poles and zeros interlace near the unit circle so as to obey the phase constraint. (Note that even the simple filter z^{-1} , corresponding to a unit sample delay, reaches phase π at half the sampling rate and is therefore not a passive transmission filter.) Physically, pole-zero interlacing corresponds to the fact that a bridge impedance “looks like a spring” at frequencies from 0 to the first resonance frequency, then it looks like a mass up to the next resonance, then like a spring again, and so on, up to half the sampling rate. These are the classical “stiffness controlled” and “mass controlled” frequency regions of a lightly damped impedance. Right on a resonance frequency, the phase goes to 0 and the impedance “looks like a dashpot” in that the impedance is real.

Note that a yielding bridge introduces losses into all attached strings. Therefore, in a maximally simplified implementation, all string loop filters may be eliminated, resulting in only one filter—the transmission filter—serving to provide all losses in a coupled-string simulation. If that transmission filter is multiply free, then so is the entire multi-string simulation.

6 Summary

Techniques applicable to efficient synthesis of stringed musical instruments were presented, along with some further extensions. Specific techniques included lumping of distributed losses and dispersion, convolving body resonators and string excitation signals into aggregate excitation look-up tables, bowed strings as periodically plucked strings, single-filter coupled strings implementation, and ways to eliminate multiplications. Since multiplies are intrinsically more expensive than additions in linear number systems (e.g., a 16 by 16 multiply requires 16 extended-precision additions), the number of voices possible in a VLSI implementation normally goes up as the number of multiplications goes down.

7 Appendix

To introduce losses into the wave equation, odd-order time derivatives such as \dot{y} , $\partial^3 y / \partial t^3$, and $\partial^5 y / \partial t^5$ are introduced. To introduce dispersion, e.g., for stiff strings and bars, a fourth-order term proportional to y'''' is added in. A general, linear, time-invariant, differential equation which covers all of these cases is

$$\sum_{k=0}^{\infty} \alpha_k \frac{\partial^k y(t, x)}{\partial t^k} = \sum_{l=0}^{\infty} \beta_l \frac{\partial^l y(t, x)}{\partial x^l}$$

On setting $y(t, x) = e^{st+vx}$, (or taking the 2D Laplace transform with zero initial conditions), we obtain the algebraic equation,

$$\sum_{k=0}^{\infty} \alpha_k s^k = \sum_{l=0}^{\infty} \beta_l v^l$$

Solving for v in terms of s is straightforward in the case of simple losses and stiff strings, and doing so yields the filtering needed to simulate simple losses and dispersion [Smith 1992]. More general cases are not solvable in closed form, but are solvable numerically. For example, note that starting at $s = 0$, we normally also have $v = 0$ (corresponding to the absence of static deformation in the medium). Stepping s forward by a small differential $j\Delta\omega$, the left-hand side can be approximated by $\alpha_0 + \alpha_1\Delta\omega$. Requiring the generalized wave velocity $s/v(s)$ to be continuous, a physically reasonable assumption, the right-hand side can be approximated by $\beta_0 + \beta_1\Delta v$, and the solution is easy. As s steps forward, higher order terms become important one by one on both sides of the equation. Each new term in v spawns a new solution for v in terms of s , since the order of the polynomial in v is incremented. It appears possible that *homotopy continuation methods* [Morgan 1987] can be used to keep track of the branching solutions of v as a function of s . For each solution $v(s)$, let $v_r(\omega)$ denote the real part of $v(j\omega)$ and let $v_i(\omega)$ denote the imaginary part. Then the eigensolution family can be seen in the form $\exp\{j\omega t \pm v(j\omega)x\} = \exp\{\pm v_r(\omega)x\} \cdot \exp\{j\omega(t \pm v_i(\omega)x/\omega)\}$. Defining $c(\omega) \triangleq \omega/v_i(\omega)$, and sampling according to $x \rightarrow x_m \triangleq mX$ and $t \rightarrow t_n \triangleq nT(\omega)$, with $X \triangleq c(\omega)T(\omega)$ (the spatial sampling period is taken to be frequency invariant, while the temporal sampling interval is modulated versus frequency using allpass filters), the left- and right-going sampled eigensolutions become

$$\begin{aligned} e^{j\omega t_n \pm v(j\omega)x_m} &= e^{\pm v_r(\omega)x_m} \cdot e^{j\omega(t_n \pm x_m/c(\omega))} \\ &= G^m(\omega) \cdot e^{j\omega(n \pm m)T(\omega)} \end{aligned}$$

where $G(\omega) \triangleq e^{\pm v_r(\omega)X}$. Thus, a completely general map of v versus s , corresponding to a partial differential equation of any order, can be translated, in principle, into an accurate, local, linear, time-invariant, discrete-time simulation. The boundary conditions and initial state determine the initial mixture of the various solution branches.

References

- [Cook 1990] Cook, P. R. 1990 (Dec.). *Identification of Control Parameters in an Articulatory Vocal Tract Model, with Applications to the Synthesis of Singing*. Ph.D. thesis, Elec. Engineering Dept., Stanford University (CCRMA). (CCRMA thesis).
- [Horner *et al.* 1993] Horner, A., J. Beauchamp, and L. Haken. 1993. "Methods for Multiple Wavetable Synthesis of Musical Instrument Tones." *Journal of the Audio Engineering Society*, 41(5):336–356.
- [Jaffe and Smith 1983] Jaffe, D. A., and J. O. Smith. 1983. "Extensions of the Karplus-Strong Plucked String Algorithm." *Computer Music Journal*, 7(2):56–69. reprinted in [Roads 1989].
- [Karjalainen *et al.* 1993] Karjalainen, M., J. Backman, and J. Pölkki. 1993. "Analysis, Modeling, and Real-Time Sound Synthesis of the Kantele, A Traditional Finnish String Instrument." *Pages 229–232 of: Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, Minneapolis*. New York: IEEE Press.
- [Karplus and Strong 1983] Karplus, K., and A. Strong. 1983. "Digital Synthesis of Plucked String and Drum Timbres." *Computer Music Journal*, 7(2):43–55.
- [Markel and Gray 1976] Markel, J. D., and A. H. Gray. 1976. *Linear Prediction of Speech*. New York: Springer Verlag.
- [Morgan 1987] Morgan, A. 1987. *Solving Polynomial Systems Using Continuation for Engineering and Scientific Computations*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- [Morse and Ingard 1968] Morse, P. M., and K. U. Ingard. 1968. *Theoretical Acoustics*. New York: McGraw-Hill.
- [Rabiner and Gold 1975] Rabiner, L. R., and B. Gold. 1975. *Theory and Application of Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- [Roads 1989] Roads, C. (ed). 1989. *The Music Machine*. Cambridge, MA: MIT Press.
- [Smith 1985] Smith, J. O. 1985. "Introduction to Digital Filter Theory." *In: Strawn, J. (ed), Digital Audio Signal Processing: An Anthology*. Los Altos, California: William Kaufmann, Inc. a shortened version appears in [Roads 1989]. Also available as Stanford University Department of Music Technical Report STAN-M-20, April 1985.

"Efficient Synthesis of Stringed Musical Instruments," J.O. Smith, Proc. Int. Computer Music Conf. (ICMC-93), pp. 64–71, Tokyo, 1993, Computer Music Association.

- [Smith 1987] Smith, J. O. 1987. *Music Applications of Digital Waveguides*. Tech. rept. STAN-M-39. CCRMA, Music Department, Stanford University. a compendium containing four related papers and presentation overheads on digital waveguide reverberation, synthesis, and filtering. CCRMA technical reports can be ordered by calling (415)723-4971 or by sending an email request to info@ccrma.stanford.edu.
- [Smith 1992] Smith, J. O. 1992. "Physical Modeling Using Digital Waveguides." *Computer Music Journal*, 16(4):74–91. special issue: Physical Modeling of Musical Instruments, Part I. Available online at <http://www-ccrma.stanford.edu/~jos/pmudw/>.
- [Smith and Gossett 1984] Smith, J. O., and P. Gossett. 1984. "A Flexible Sampling-Rate Conversion Method." *Pages 19.4.1–19.4.2 of: Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, San Diego*, vol. 2. New York: IEEE Press. an expanded tutorial based on this paper and associated free software are available online at The Digital Audio Resampling Home Page: <http://www-ccrma.stanford.edu/~jos/resample/>.
- [Van Duyne and Smith 1992] Van Duyne, S. A., and J. O. Smith. 1992. "Implementation of a Variable Pick-Up Point on a Waveguide String Model with FM/AM Applications." *Pages 154–157 of: Proceedings of the 1992 International Computer Music Conference, San Jose*. Computer Music Association.
- [Van Valkenburg 1960] Van Valkenburg, M. E. 1960. *Introduction to Modern Network Synthesis*. New York: John Wiley and Sons, Inc.
- [Weinreich 1977] Weinreich, G. 1977. "Coupled Piano Strings." *Journal of the Acoustical Society of America*, 62(6):1474–1484. also see *Scientific American*, vol. 240, p. 94, 1979.