Developments for the Commuted Piano

Scott A. Van Duyne
CCRMA, Stanford University
savd@ccrma.stanford.edu

Julius O. Smith III
CCRMA, Stanford University
jos@ccrma.stanford.edu

ABSTRACT: We present here three developments for the Commuted Piano Synthesis model described more fully elswhere in these proceedings [3]: (1)a theoretical foundation and calibration scheme for the required linearized piano hammer system; (2) a simple algorithmic synthesis approach for the commuted soundboard impulse response, eliminating the need for any wave table memory; and (3) a calibration method for the coupled string system, required for high quality two-stage piano tone decay.

1 Background

Much study has been made of the piano and its parts, with an eye toward better understanding of the acoustics, toward more reliable numerical modeling of the piano physics, toward the development of high quality sound synthesis, and toward the development of cost-effective sound synthesis. Our interest is in the catagory of sound synthesis informed by physics and acoustics. A high quality physical approach to piano synthesis has been suggested which combines the Wave Digital Hammer [5] with coupled string synthesis [2], using the 2D Digital Waveguide Mesh [7] as a soundboard resonator, and using allpass filtering methods to stiffen the soundboard and the piano strings [6]. On the other hand, we have proposed an algorithmic method of synthesizing the inharmonic piano tones constructing spectral regions with specially tuned FM oscillator pairs [4]. More recently, we have developed a hybrid method known as Commuted Piano Synthesis [2, 3] which combines the high quality and control of physical modeling synthesis with the cost-effectiveness of sampling "synthesis". This method takes advantage of the commutativity of linear systems, and replaces the high order soundboard resonator with its own sampled impulse response played into the string at its excitation point. In order to implement Commuted Synthesis, we must linearize the hammer-string interaction, which is the focus of the first part of this paper. We have further hybridized and simplified the Commuted Piano Synthesis method by replacing the sampled soundboard impulse response with a simple algorithmic synthesis method which idealizes the soundboard and makes its physical charateristics easier to control. Lastly, presented here, is a method to calibrate the coupled piano string algorithm to real physical data using only a simple recording of a hammer hitting one string.

2 Linearizing the Piano Hammer

A fully physical nonlinear model of the hammer-string system has been offered already [5]. However, in order to implement Commuted Piano Synthesis [3], we must commute the resonant soundboard system through the hammer system to the point of excitation in the commuted piano model. This requires that we replace the entire hammer-string interaction with a linear filter. Rather inconveniently, the hammer-string interaction is highly /it nonlinear in two important respects: First, the felt itself is nonlinear in that it gets stiffer the more it is compressed. Second, the hammer leaves the string at some point, which corresponds to a shift in the models from a string interacting with a hammer to a string vibrating freely.

2.1 Linearized Analysis of the Piano Hammer-String System

Impedance of the Un-Terminated Ideal String The impedance experienced at some point on an *un-terminated* string is purely resistive:

$$R_S \triangleq F_s/V = 2R_0, \tag{1}$$

where F_S and V are the Laplace transforms of force and velocity at the driving point and R_0 is the wave impedance of the string, which is dependent on the squre root of string tension times string density. The $2R_0$ in the above equation results from taking into account the impedance of both halves of the string, as seen at the driving point.

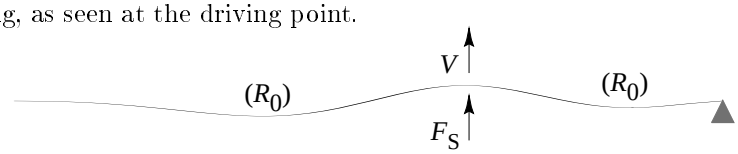

Figure 1: *String Terminated on One Side Only*

Impedance of the Terminated Ideal String In the case of the piano hammer-string interaction, waves from the agraffe return and interact with the hammer before it leaves the string for most notes. However, the return waves from the bridge end of the string do not make it back before the hammer leaves the string, except in the very highest notes. Therefore, we formulate a half terminated string impedance taking into account a one sided termination at the agraffe end, as shown in Figure 1. The velocity response of a force impulse at the strike position is an impulse followed by an inverted impulse which returns reflected off the essentially rigid agraffe end of the string T seconds later:

$$V = \frac{F_s}{2R_0}\left(1 - e^{-sT}\right) \quad \Longrightarrow \quad R_S \triangleq \frac{F_s}{V} = \frac{2R_0}{1 - e^{-sT}} \tag{2}$$

Impedance of the Ideal Linear Hammer Let us assume that the hammer is of the form shown in Figure 2, essentially a mass and spring system, where the spring represents the felt portion of the hammer. We find that the impedance relation is:

$$F_H = R_H\left(V - \frac{v_0}{s}\right) \quad \text{where} \quad R_H \triangleq \frac{ks}{s^2 + k/m} \tag{3}$$

and where v_0/s represents the step input of the intial striking velocity. R_H has a zero at DC and two conjugate poles indicating an oscillation frequency of $\sqrt{k/m}$.

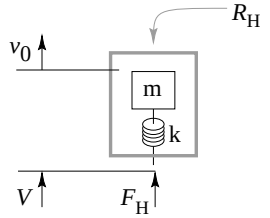

Figure 2: *The Linear Mass-Spring Hammer Model*

Connecting the Hammer to the String When the hammer is in contact with the string, we take the velocity of the string equal to the velocity of the spring end of the hammer, and the force on the string equal and opposite to the force on the spring, $F_S = -F_H$. Plugging in the string impedance relation, $V = F_s/R_s$, we find:

$$F_S = -F_H \triangleq -R_H\left(V - \frac{v_0}{s}\right) = -R_H\left(\frac{F_S}{R_S} - \frac{v_0}{s}\right) \tag{4}$$

$$\Longrightarrow F_S = \left(\frac{R_H R_S}{R_H + R_S}\right)\frac{v_0}{s} \triangleq \left(R_H \Big\| R_S\right)\frac{v_0}{s} \tag{5}$$

In the *unterminated* string case, we define H_∞ as the transfer function from the initial striking velocity step to the force experienced by the string (and, equivalently, by the hammer felt). Taking the hammer to be a simple mass-spring system, we find that the H_∞ transfer function is now a damped second order system, which looks just like the R_H except for the under bracketed damping term (6). For practical physical parameters, this is an over damped system with real poles.

$$H_\infty \triangleq R_H \Big\| R_S = \left(\frac{ks}{s^2 + k/m}\right) \Big\| 2R_0 = \frac{ks}{s^2 + \underbrace{\frac{k}{2R_0}s}_{} + \frac{k}{m}} \tag{6}$$

For the one side terminated string case, we define H_T. Again, we find H_T is like R_H but for the under bracketed damping term, which in this case contains an interesting time delay part.

$$H_T \triangleq R_H \Big\| R_S = \left(\frac{ks}{s^2 + k/m}\right) \Big\| \left(\frac{2R_0}{1 - e^{-sT}}\right) = \frac{ks}{s^2 + \underbrace{\frac{k}{2R_0}\left(1 - e^{-sT}\right)s}_{} + \frac{k}{m}} \tag{7}$$

2.2 Implementation

Conveniently, we find a recursion relationship between H_∞ and H_T, which is independent of the exact nature of the hammer impedance, R_H.

$$H_T = \frac{H_\infty}{1 - e^{-sT}H_\infty} \tag{8}$$

This allows a simple recursive hammer filter implementation of the form in Figure 3:

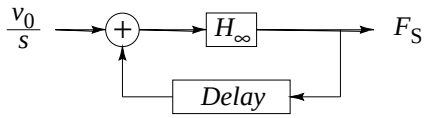

Figure 3: *Step-Driven Recursive Hammer Filter*

Since, in this case, the hammer never leaves the string (from the linear system assumption), we may prefer to include a cutoff envelope in the feedback loop to terminate the reflections from the agraffe at some point, or better, break out the first few reflections in a feed forward formulation as in Figure 4:

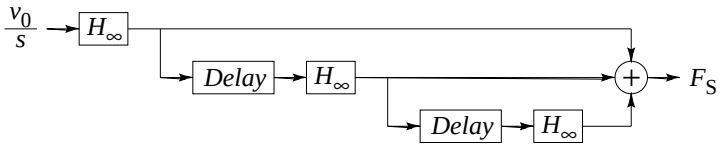

Figure 4: *Feed Forward Hammer Filter*

Noting that H_∞ is a differentiated lowpass filter,

$$H_\infty = \frac{ks}{s^2 + \frac{k}{2R_0}s + \frac{k}{m}} = sL_p \tag{9}$$

the step-driven hammer system of Figure 3 may be commuted to an impulse-driven system, as required for Commuted Piano Synthesis [3]. This is shown in Figure 5. In this formulation, the hammer feedback loop contains what is fundamentally a lowpass filter and a DC-blocker. It is easy to break this out into a feed forward form as in Figure 4.

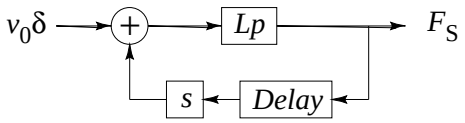

Figure 5: *Impulse-Driven Recursive Hammer Filter*

2.3 Analysis of Real Hammer-String Interaction Data

Using the Wave Digital Hammer [5] parameterized with measured data provided by [1], we were able compute the forces experienced by terminated and un-terminated middle-C strings during a hard hammer strike. In the upper left plot of Figure 6, we see the felt compression force curves for a hammer hitting an unterminated middle-C string (dashed line) and a terminated middle-C string (solid line). The multiple pulses correspond to return waves from the agraffe interacting with the hammer while it is still in contact with the string. Note that in the unterminated string case, the force curve ramps smoothly to zero and the hammer apparently comes to rest on the string as in an over damped second order system. However, in the terminated string case, the hammer leaves the string when the return waves finally throw it away. The upper right plot of Figure 6 show the dB magnitude spectra of these force curves. Note, here, that the overall bandwidth of both the terminated string and unterminated string hamme shock spectra are about the same. The multi-pulse spectrum (solid line) differs from the single-pulse spectrum (dashed line) primarily in a slight ringing of the lower spectrum region.

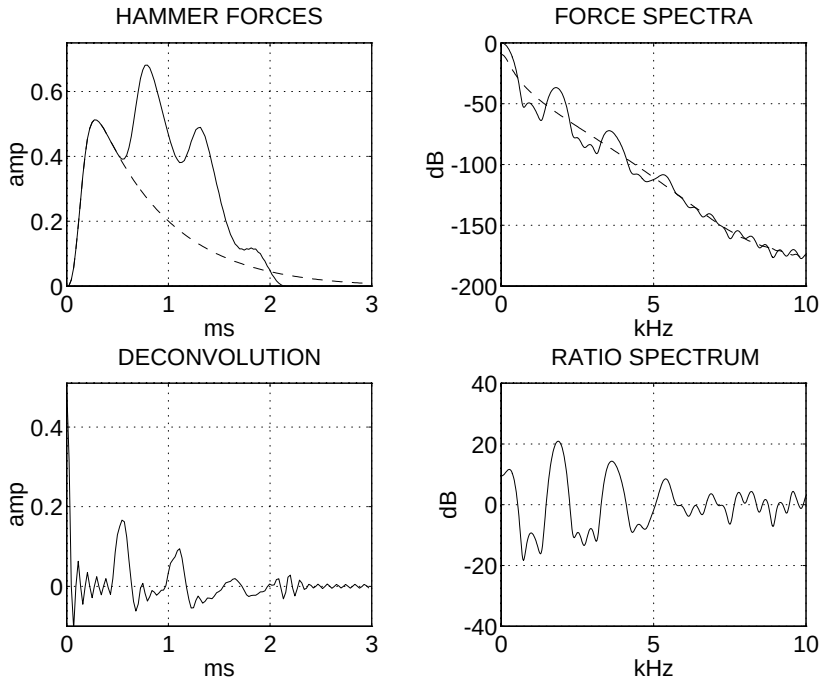

Figure 6: *Middle-C Struck Hard: Force signals computed using the WDH parameterized with physical data taken from Chainge and Askenfelt (1994)*

The lower right plot is the dB magnitude of the complex ratio of the multi-pulse spectrum and the single-pulse spectrum. The several low frequency lobes correspong to the spectral peaks one would expect from the hammer staying in contact with the string at the stike point (about 1/8 the way along the string) for some finite duration. It is of some interest that keeping the hammer in contact with the string introduces spectral peaks about every eight harmonics, whereas an impulsive

strike at the same position on the string introduces spectral *nulls* every eight harmonics. The piano hammer interaction is a compromise between these two extremes of behavior.

We further note that there are odd looking wiggles in the ratio spectrum, clearly visible around the 5–10 kHz range. These correspond in width to the side lobes one would expect from rectangularly windowing the time domain signal at exactly the point where the hammer leaves the terminated string. Hence, the severe nonlinear effect of the hammer leaving the string (which changes the entire linear system model) turns out in the spectral domain to be a simple convolution by the appropriate rectangular window sinc function.

The lower left plot in Figure 6 shows the inverse transform of the ratio spectrum. This is what is left of the multi-pulse hammer force signal if we de-convolve the single pulse force signal out of it. It appears to be a recursively damped impulse train, with some DC blocking, eventually centening the signal around zero. This is what was predicted by the linear hammer analysis as shown in Figure 5.

2.4 Spectral Modeling Approach to the Multi-Pulse Effect

An alternative approach to hammer filter design is to model the complex ratio spectrum as shown in the lower right plot of Figure 6 directly as a low order filter. This reduces the recursive or feed forward filter design methods of modeling multiple force pulses to a simple spectral equalization filter, as shown in Figure 7.

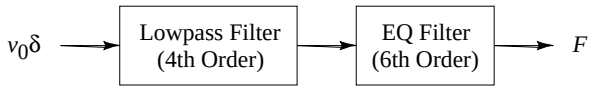

Figure 7: *Spectral EQ Method of Modeling the Hammer Filter*

In this case, we used a fourth order fit for the single pulse hammer lowpass filter. We then made a sixth order equalization filter fit to a few of the significant low frequency features of the ratio spectrum. The right hand plot of Figure 8 shows the equalization filter fit. The left hand plot shows the time domain output of the hammer filter system shown in Figure 7. The thick dotted lines are actual data as generated by the Wave Digital Hammer and the solid lines are the result of the filter fits. Note that the phase information in the sixth-order ratio spectrum fit results in a very good time domain approximation. In general, the coefficients of the lowpass filter part of this structure will be highly dependent on strike velocity, the harder the strike, the wider the bandwidth. However, the equalization part of this structure is reasonably consistant over strike velocity, and, in the simplified model, may be held constant over strike velocity, although it will vary over piano key number.

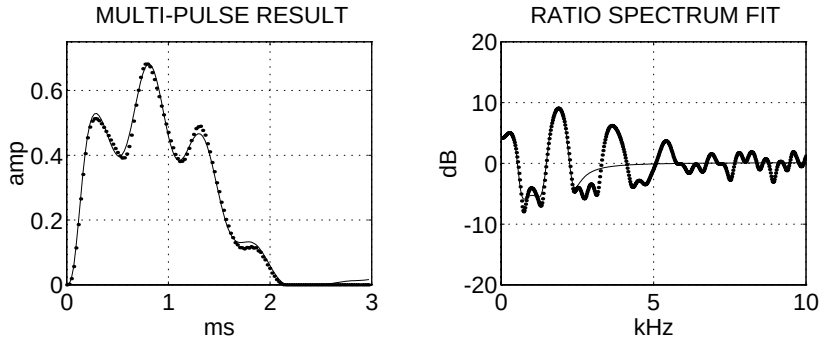

Figure 8: *Sixth Order Filter Fit to Ratio Spectrum*

3 Excitation Synthesis with Nonlinearly Filtered Noise

The impulse response of the piano soundboard is fundamentally a superposition of many exponentially decaying sinusoids, at least in its linear approximation. The reverberant effect of the soundboard occurs as energy from the struck piano strings is coupled into these modes and reverberated. However, if there were some particular modes of the soundboard which were unusually prominent, exhibiting a clear peak in the impulse response spectrum, and having an unusually long decay time, then a string which contained this frequency in one of its partials would couple into this mode more significantly than a string which did not have that frequency among its partials. This could produce unwanted unevenness in the piano tone from note to note. In general, much effort has been applied to the design of real piano soundboards to avoid such situations as these. The idealized piano soundboard should have a smooth, or flat spectral response locally, although it is evident that higher frequency modes decay a little faster that low frequency modes.

It is difficult to design a resonant system with such a flat response without using a very high order filter, for example the 2D Digital Waveguide Mesh [7]. On the other hand, it is easy to model the impulse response of such a system as exponentially decaying white noise, with the possible extension of a time varying low pass filter applied to model high frequency modes decaying more quickly than low frequency modes. Using such a nonlinearly filtered noise model, we may synthesize any number of reverberant systems which have the characteristic that they have more or less smooth responses over the frequency spectrum, with no particular peaks of importance. The piano soundboard is a system of this kind.

In Figure 9 we show such a soundboard impulse response synthesis system. White noise is being fed into a time varying low pass filter whose gain and bandwidth are both being controlled by envelopes. One possible implementation of this would use a one-pole low pass filter whose denominator coefficient is being swept toward -1, thereby shrinking the bandwidth. If the numerator coefficient is modified to keep gain at DC constant, the amplitude envelope might even be dispensed with in a simplified system. Alternatively, more elaborate noise filtering systems may be used, possibly breaking the noise into frequency bands which would be enveloped independently to calibrate to some particular impulse response.

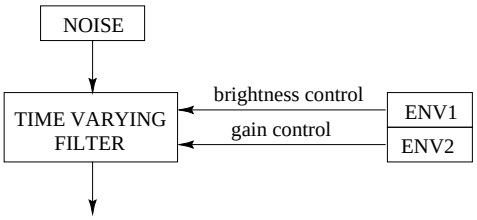

Figure 9: *Synthesis of Soundboard Tap with Nonlinearly Filtered Noise*

Synthesizing Sustain Pedal Effect Just as the dry soundboard impulse response may be commuted to the point of excitation, similarly we may commute the entire sampled impulse response of the soundboard plus open strings with dampers raised to the point of excitation to obtain the resonant effect of the sustain pedal. Further, since there are so many resonating partials, the spectral response is essentially flat and filtered white noise with a long slow decay rate makes a good synthetic approximation.

4 Calibrating Coupled Stings

4.1 The Coupled String Model

Figure 10 illustrates a coupled piano string model for one note of the piano. The Coupling Filter represents the loss at the yielding bridge termination, and controls the coupling of energy between

and among the three strings. Each of the three string loops shown contain two Delay elements, the first corresponding to the delay path from the hammer strike point to the agraffe and back, the second corresponding to the delay path from the hammer strike point to the bridge and back. The relative delay length ratio for most strings is about 1 to 8, although the relative delay lengths may be set to model any particular piano string strike position. The input signals E_1, E_2, and E_3 are taken from the output of the hammer filter, which has been driven, in turn, by a soundboard impulse response, or a nonlinearly filtered noise excitation synthesis. Note that the input signals are introduced into the string loops at two points, in positive and negative form: this models the spectral combing effect of the relative strike position of the hammer on the string.

The signals C_1, C_2, and C_3 should be set to 1.0 during the sustain portion of the piano sound, and should be ramped to some appropriate loop attenuation factor, such as .95, at key release time. Alternatively, some more elaborate release sound model might be used. Note that, for *una corda* pedal effects, one or more of the signals E_1, E_2, or E_3, should be set to zero at key strike time. This causes the coupled string system to move quickly into its second stage decay rate, just as is found in real piano sounds when the *una corda* pedal is depressed.

In this coupled string model, the delay lengths are fine-tuned such that the effective pitch of each of the three string loops is *very nearly* equal, but *not exactly* equal. This is the mechanism by which two stage decay is synthesized in the commuted piano synthesis model. The Stiffness Filters, as shown in the Figure, are intended to be an allpass filter structure which modifies the phase response of the loop so as to create the effect of the natural inharmonicity of the piano string partials. We recommend a bank of one-pole allpass filters as described in [6].

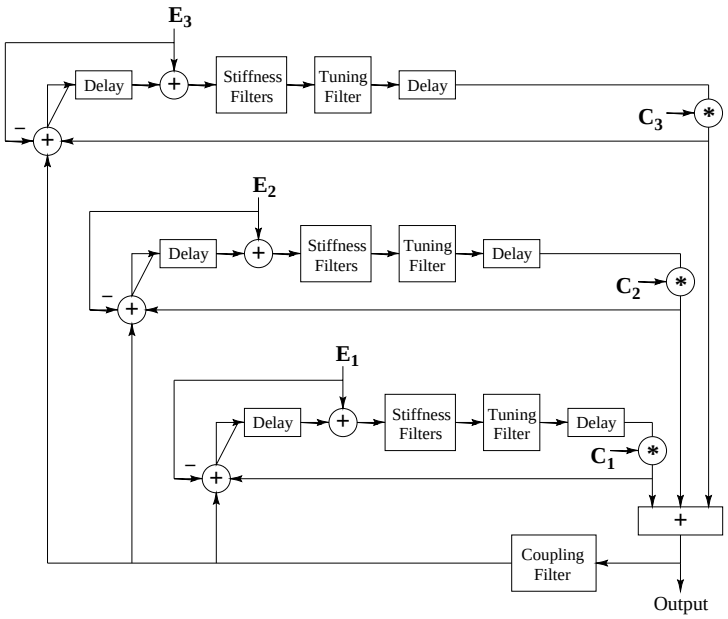

Figure 10: *Three Piano Strings Coupled at a Bridge Termination*

4.2 Calibrating the Coupling Filter

Ideally, from a physical perspective, we would like to measure empirically the bridge impedance, R_b, and the string impedance, R_0; and then from these measurements compute the desired Coupling Filter. However, following the spirit of the simplified string loop model presented above, let us say we have already calibrated a *single* string system and know $LP(z)$, a lowpass filter modeling the per period attenuation of the tone, and $AP(z)$, an allpass filter summarizing the dispersion in the string due to stiffness. We have presumably done this by measuring the partial frequencies and

corresponding decay rates of a single piano string. This may be accomplished by physically damping two of the three piano strings in a piano note group with felt, rubber, or some such means, and then recording the sound of the remaining undamped string decaying after it is struck. The decay rate of this single string should not contain very much two-stage decay interference from the other damped strings, but should, instead, produce a reasonable single stage decay from which data about the partial frequencies and their individual decay rates may be extracted.

Loss in a string-bridge system comes almost entirely from the bridge termination itself. That is, the loss from viscous air drag and internal friction is very small compared to termination loss. Therefore, let us simply say that $LP(z) \triangleq T_f(z)$ is the force wave transfer function at the bridge, that the string is rigidly terminated at the other end so that the force wave transfer function there is unity, and that the dispersion, $AP(z)$, is entirely due to stiffness in the string, and not due to any significant reactive qualities in the bridge. We may therefore write [2],

$$LP(z) \triangleq T_f = \frac{R_b - R_0}{R_b + R_0} \tag{10}$$

and solve for R_b in terms of LP,

$$R_b = R_0 \frac{1 + LP(z)}{1 - LP(z)} \tag{11}$$

The coupling filter for N strings coupled at an impedance R_b is [2]

$$H_b \triangleq \frac{2}{N + R_b/R_0} = \frac{2}{N + \frac{1 + LP(z)}{1 - LP(z)}} = \frac{2(1 - LP)}{(1 + N) + (1 - N)LP} \tag{12}$$

In summary of this calibration approach, we have measured the sound of a single string decaying, derived the loop loss filter from this data, then taken this to be the force wave transfer function at the bridge (since we assume that most all of the loss is due to yielding bridge, and the internal string loss is small in this situation); from this point, we derive the bridge impedance and thence the N-string coupling filter.

In the model shown in Figure 10, we have three strings coupled, $N = 3$. However, several minus signs have been commuted around in that figure and the Coupling Filter is actually represented by $-H_b$. To complete the model, the Tuning Filters should be tweaked by a good piano tuner to achieve a fine, full-bodied two-stage decay rate (around 0.5–2 Hz detuning between strings).

References

[1] Chaigne, A. and A. Askenfelt. "Numerical simulations of piano strings." I & II. *JASA* **95** (2) and (3). 1994.

[2] Smith, J. O. "Efficient Synthesis of Stringed Musical Instruments." *Proc. ICMC*, Tokyo. 1993.

[3] Smith, J. O. and S. A. Van Duyne "Commuted Piano Synthesis." *Elsewhere in these Proceedings.*

[4] Van Duyne, S. A. "Low Piano Tones: Modeling Nearly Harmonic Spectra with Regions of FM." *Proc. ICMC*, San Jose. 1992.

[5] Van Duyne, S. A.; Pierce, J. R. and J. O. Smith. "Traveling Wave Implementation of a Lossless Mode-Coupling Filter and the Wave Digital Hammer." *Proc. ICMC*, Århus. 1994.

[6] Van Duyne, S. A. and J. O. Smith. "A Simplified Approach to Modeling Dispersion Caused by Stiffness in Strings and Plates." *Proc. ICMC*, Århus. 1994.

[7] Van Duyne, S. A. and J. O. Smith. "Physical Modeling with the 2-D Digital Waveguide Mesh." *Proc. ICMC*, Tokyo. 1993.