

Enhanced Wave-Based Modelling of Musical Strings. Part 1: Plucked Strings

Hossein Mansour¹⁾, Jim Woodhouse²⁾, Gary P. Scavone¹⁾

¹⁾ Computational Acoustic Modeling Laboratory, Schulich School of Music, McGill University,
555 Sherbrooke Street West, Montréal, Québec H3A 1E3, Canada

²⁾ Cambridge University Engineering Department, Trumpington Street, Cambridge CB2 1PZ, UK.
jw12@cam.ac.uk

Summary

A physically-accurate time-domain model for a plucked musical string is developed. The model incorporates detailed dispersion and damping behaviour measured from cello strings, and a detailed description of body response measured from a cello body. The resulting model is validated against measured pizzicato notes using the same strings and cello, and good accuracy is demonstrated. The model is developed in a form that makes extension to the case of a bowed string very straightforward.

PACS no. 43.40.Cw, 43.75.Gh

1. Introduction

This paper presents a refined simulation model of the motion of a plucked string, with a focus on achieving high physical accuracy by incorporating the most complete theory and measurement data available. Since this model draws upon best practice from earlier research, the description involves an element of review. However, significant new measurements and validation experiments are also included. In an earlier study, several methods for accurate synthesis of guitar plucks were compared [1]. The best performance was obtained using a frequency-domain approach, but for the purposes of musical synthesis a time-domain approach is preferable because of the latency implicit in the frequency-domain method. A time-domain travelling-wave approach was also tried in [1], but was found to perform relatively poorly. One aim of the present work is to improve the implementation of this model and demonstrate that it can work well.

The model is developed in such a way that it can also be used for bowed strings, and this is another strong motivation for needing a time-domain methodology: the nonlinear friction force in a bowed string can only be handled in the time domain, if transient simulations are wanted. As a consequence, parts of the model are developed in a form that is slightly more complicated than would be needed for plucked strings alone. Also, most of the detailed results to be presented here concern the cello. Calibration measurements on cello strings and a particular cello body will be used to illustrate the approach, and comparisons will

then be shown between synthesised and measured pizzicato notes on that cello. The application of the model to bowed string motion is described in a companion paper [2].

A primary goal is to make the model physically accurate and to keep the link between the model and physical parameters as clear as possible. This contrasts with the priorities in the sound synthesis field, where physical details may be compromised to improve computational efficiency as long as their exclusion does not significantly worsen the quality of the synthesised sound. Having said that, the two fields have remained closely knit: indeed, the methods used here to model the damping and dispersion of a string are tailored versions of models originally developed for sound synthesis purposes.

There is a long history of theoretical analysis of vibrating strings [3]. In 1746, d'Alembert [4] published a solution for the motion of an ideal lossless string in the form of a general superposition of two waves travelling in opposite directions with speed $c_0 = \sqrt{T_0/m_s}$, where T_0 is the string's tension and m_s is its mass per unit length. Much more recently, this idea formed the basis of a successful modelling strategy for a bowed string [5], [6], which evolved into what has become known as “digital waveguide modelling” (see for example Smith [7]). This is the approach followed in the present work.

When applied to a plucked string, the method is very simple. The assumed details of any particular pluck can be used to determine the initial shapes of the waves that travel in the two directions. A pluck involves initial application of a force at a particular point on the string (or over a short length of string), this force jumping to zero at the moment of release of the string. This contrasts with the situation in a bowed string, where force is continuously

Received 15 June 2016,
accepted 12 September 2016.

applied through the bow hairs to the string. In that case, the incoming waves at the bowed point interact with the friction force at the bow to generate outgoing waves (see for example [6]). For the plucked-string case there is no force at the plucking position, so the waves simply cross at this point to become unaltered outgoing waves. Linear theory is assumed throughout this work, and so the incoming waves returning to the pluck/bow position at any given time step in the simulation process can be calculated by convolution of the outgoing waves at earlier times with suitable convolution kernels.

The process of modelling consists essentially of determining these kernel functions in order to represent the relevant physical processes to sufficient accuracy. The two kernels are traditionally called “reflection functions”, denoted h_1 and h_2 for the bridge and finger sides respectively (“finger” is used as a shorthand for finger/nut throughout). In order for h_1 and h_2 to be physically accurate, they must satisfy

$$\int_{-\infty}^{\infty} h_1 dt = \int_{-\infty}^{\infty} h_2 dt = -1. \quad (1)$$

If this condition is not met, the mean values of the left- and right-going travelling waves can drift, which in physical terms would correspond to the entire string shifting position.

For a perfectly flexible and lossless string with rigid terminations, both reflection functions consist simply of delayed and inverted unit delta functions. The required delay to produce a desired fundamental frequency f_0 for the complete string is equal to β/f_0 for the bridge side function h_1 and $(1-\beta)/f_0$ for the finger side function h_2 , where β is the distance of the excitation point from the bridge, expressed as a fraction of the total string length. A more realistic model requires more complicated reflection functions, but traces of this simple structure will remain in evidence.

2. Model ingredients and implementation

There are several aspects of underlying physics relevant to a plucked string. Some are intrinsic to the string itself, determining the details of dissipation and dispersion. Others involve coupling to the vibration modes of the instrument body, which also induces coupling between the two polarisations of string motion. At the other end of the vibrating string, the player’s finger and the details of contact with a fingerboard or fret may have an influence. Finally, there are features of a complete musical instrument that might influence a given plucked or bowed note: the vibration of non-excited sympathetic strings, and the vibration of the after-lengths of the strings on the far side of the bridge, including their interaction with the tailpiece. All these factors can be included in the model to be presented here.

2.1. Dispersion and dissipation in the string

2.1.1. Theoretical background

All real strings exhibit non-zero bending stiffness and frequency-dependent dissipation. In much of the earlier

work on plucked and bowed strings (see for example [8, 9, 1]) these factors were represented via approximate analytic reflection functions, but more sophisticated representations based directly on measurements will be developed here. The approach is implemented in the time domain, but the reflection functions can be designed to match frequency-domain characteristics: in other words, they can be viewed as the impulse responses of filters with particular magnitude and phase characteristics. This will allow the use of modern digital filter design methods. Following the convention of the musical synthesis literature, these will be called “loop filters” throughout.

The standard equation for the free motion of a stiff string without damping is

$$EI \frac{\partial^4 y}{\partial x^4} - T_0 \frac{\partial^2 y}{\partial x^2} + m_s \frac{\partial^2 y}{\partial t^2} = 0, \quad (2)$$

where for a solid string E is the Young’s modulus and I the second moment of area of the string’s cross-section. For a typical layered musical string, the combined parameter EI is best regarded as an empirical factor, to be determined by measurement. The mode shapes remain very similar to those of a perfectly flexible string, but the natural frequencies are no longer exactly harmonic. The bending stiffness produces a wave propagation speed that is frequency dependent, which results in a “stretching” of the natural frequencies. Rayleigh’s principle can be used to show that the n th natural frequency of a stiff string is given by

$$f_n \approx n f_0 \sqrt{1 + B n^2} \approx n f_0 \left(1 + \frac{B n^2}{2} \right), \quad (3)$$

where f_0 is the first mode frequency if the string had been perfectly flexible, and the inharmonicity coefficient B is given by

$$B = \frac{EI \pi^2}{T_0 L^2}, \quad (4)$$

where L is the length of the string.

The inharmonicity of many musical strings is known to be above the threshold for human perception [10, 11], so it can be of direct perceptual significance. The systematic stretching revealed in Eq. (3) also results in the pitch being perceived slightly sharper than the frequency of the fundamental. A degree of inharmonicity is essential to the normal sound of some instruments, such as the modern piano [12, 13], but too much of it is certainly not desirable. A familiar way to limit the inharmonicity of low-frequency strings in practice is to use a thin core over-wound with one or more layers of wire to give the desired mass per unit length without adding too much to the bending stiffness EI .

It should be noted that the fourth-order equation of motion, Equation (2), results in four solutions, only two of which are naturally included in the travelling wave approach; the other two are a pair of fast-decaying quasi-evanescent waves. These waves are only important in the vicinity of the excitation point, and within a short period

of time after the excitation. Ducasse has estimated those limits for a piano C_2 string to be in the neighborhood of 2 cm and within 0.1 ms of the hammer excitation [14]. For thinner strings, like those of a cello or a violin, the spatial limit should be even smaller, but it is still of the order of the bow width and is likely to be important in the detailed interaction of a bow with a string [15]. However, these evanescent waves will be ignored in the model to be developed here.

On a stiff string the group velocity rises with increasing frequency, resulting in the formation of “precursor” waves preceding the main peak in the reflection function. An approximate expression for this reflection function was presented by Woodhouse [8] (see Figure A1), and used in subsequent work. Equation (2) becomes non-physical at very high frequencies because the wave velocity rises without limit, whereas any real material has a maximum possible wave speed. In consequence, to use the analytical expression in simulations it is necessary to filter it with some chosen cutoff frequency. A way of avoiding this requirement will be presented in Section 2.1.3.

In earlier work, string damping was also often represented by an analytic formula, in this case a rather crude one. A form of reflection function was introduced in [16] and then used in several later studies [17, 18], which attempts to give the same Q factor to all string modes. The function for the bridge side takes the form

$$h_1 = \frac{2\beta L / (2Qc_0)}{\pi [(t - 2\beta L / c_0)^2 + (2\beta L / (2Qc_0))^2]}, \quad (5)$$

while for the finger side, β must be replaced by $(1 - \beta)$. Note that a reflection function designed according to Equation (5) is symmetric around its peak which is expected as it is the impulse response of a linear-phase loop filter.

The design of reflection functions based on Equation (5), or any other FIR filter for that matter, can become problematic for short segments of lightly damped strings. The discrete-time form of such functions will have only a few significantly non-zero elements, so that normalisation of the area in order to satisfy the discrete version of Equation (1) might require a large adjustment to the peak height, and hence produce a large deviation from the desired behaviour. The problem will be illustrated in Section 3 by simulation of an open D_3 cello string using this type of reflection function, compared with the alternative formulation that will now be developed.

2.1.2. Measurements of string damping

To do better than the early models, it is first necessary to have reliable data for the intrinsic damping of the string. The damping of the first 30 modes, characterised by Q factors, was measured [19] for seven sets of nominally-identical “D’Addario Kaplan Solutions” cello strings (model KS510 4/4M). The inharmonicity coefficients were determined at the same time. The measured Q factors for each string mode were averaged across the different strings tested, to minimise the effect of manufacturing variations and experimental uncertainty. The measure-

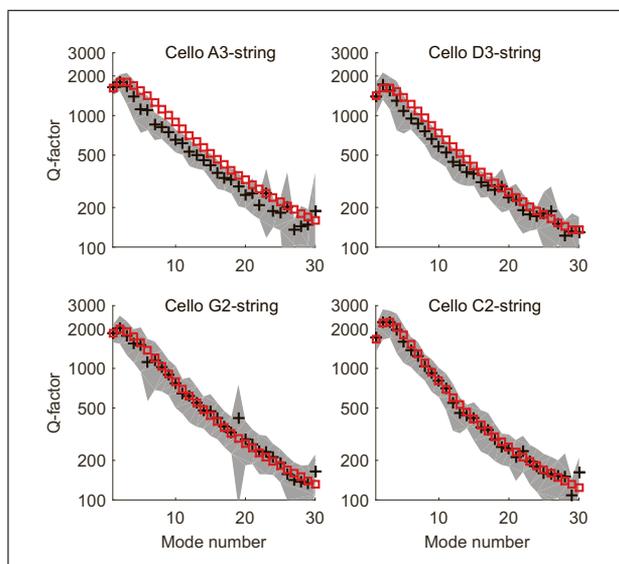


Figure 1. Average measured Q factor (plus signs) plus/minus one standard deviation (grey shade) for D’Addario Kaplan Solutions cello strings. The red squares show the fit of Equation (6) to the measured data.

ments were made on a rigid granite base so that the results only correspond to the intrinsic damping of the strings.

A model due to Valette [20] was then used to give a parametric fit to the measurements: such a fitted model allows simulation of different notes played on a given string. This model considers the net effect of viscous damping by the surrounding air, viscoelasticity and thermoelasticity of the string material, and internal friction. Viscoelasticity and thermoelasticity both manifest themselves by creating a complex Young’s modulus, which comes into the equation of motion through the bending stiffness term. Its significance increases with the square of the frequency. Aerodynamic loss predominantly affects the lower frequencies, while internal friction has a rather uniform influence on all frequencies. In mathematical form, the Q factor of the string’s n th mode is expressed as

$$Q_n = \frac{T_0 + EI(n\pi/L)^2}{T_0(\eta_F + \eta_A/\omega_n) + EI\eta_B(n\pi/L)^2}, \quad (6)$$

where ω_n is the angular frequency, and η_F , η_A and η_B are coefficients determining “friction”, “air” and “bending” damping respectively. These three coefficients can be estimated by fitting Equation (6) to the measured Q factors. Both measured and fitted data are illustrated in Figure 1; the shaded band indicates ± 1 standard deviation to show the variability of measurements. The fitted parameter values, as well as other string properties, are summarised in Table I.

The pattern of the Q factors looks almost identical across the four cello strings, when plotted against the string mode number (as opposed to the mode frequency). It can be seen in Figure 1 that Valette’s proposed relation gives a better fit to the Q factor trend of the C_2 and G_2 strings than it does to the D_3 and A_3 strings. For the D_3 and

Table I. Measured and estimated properties for a set of D'Addario Kaplan Solutions cello strings. All parameters are relevant to the transverse vibrations, and the effective length of the open strings is assumed to be 690 mm.

Tuning			A_3	D_3	G_2	C_2
Frequency	f_0	Hz	220	146.8	98	65.4
Tension	T_0	N	171	135.9	135.5	131.5
Mass/unit length	m_s	g/m	1.85	3.31	7.40	16.14
Bending stiffness	EI	$10^{-4} N/m^2$	3.26	2.48	1.88	6.20
Inharmonicity	B	10^{-6}	39.5	37.9	28.7	97.8
Characteristic impedance	Z_0	Kg/s	0.56	0.67	1.00	1.46
Loss coefficients	η_F	10^{-5}	22	23	20	12
	η_B	10^{-2}	11.4	12.5	13	4.7
	η_A	1/s	0.12	0.11	0.04	0.07

A_3 strings, the decrease of the Q factors beyond their peak value is steeper than is predicted by Valette's model. For all strings, the highest Q factor occurs at the second or the third mode, with the maximum values ranging from 1200 to 3000. This observed trend of Q factors for cello strings is significantly at odds with the ones earlier reported for harpsichord strings [20] and guitar strings [21, 11]: all these other types of musical string showed the maximum of Q factor occurring at much higher mode numbers. Presumably the pattern observed in the cello strings is a deliberate consequence of their elaborate multi-layer construction: given that construction, it is perhaps no great surprise that Valette's simple model does not quite succeed in capturing the frequency variation correctly.

A final note on the frequency-dependent Q factor concerns the case of finger-stopped strings. Stopping the string at one end by the finger will introduce significant additional damping, particularly for instruments like those of the violin family that do not have frets. In a study by Saw [22], the damping of a finger-stopped string was compared to that of an open string. Those results suggest a simple way to represent, roughly, the effect of finger damping: η_F should be tripled, while keeping η_A and η_B unchanged.

2.1.3. Filter implementation

To accurately account for the damping trend of a string over the desired range of frequencies, the reflection functions must implement the frequency-dependent attenuation factors over their corresponding string lengths. These reflection functions can be viewed as the impulse responses of frequency-domain filters that implement the desired attenuation trends. Considering the bridge side of the string, there are $\beta f/f_0$ cycles of frequency component f in a round trip to and from the bridge. Therefore, the gain G_1 of the filter for the bridge side is related to the desired Q factor by

$$G_1(f) = e^{-\pi\beta f/f_0 Q}, \quad (7)$$

directly from the definition of Q factor as π times the number of periods for the amplitude to decay by the factor $1/e$. The corresponding expression for gain G_2 for the finger side is obtained by replacing β with $(1 - \beta)$.

Damping will be implemented separately from dispersion, so the first stage is to find the loop filter for a damped

but non-dispersive string on which all frequencies travel with the same propagation speed (i.e. is linear-phase).

Using the parameters from Table I, the desired gain factor, or response magnitude, over the full range of frequencies and for each note was calculated by combining Equations (6) and (7). The DC gain was set to unity to comply with Equation (1), and for finger-stopped notes η_F was tripled. Equation (6) naturally limits the Q factor at high frequencies to the value $1/\eta_B$, around 20 for these cello strings; however, for practical reasons concerning the filter design procedure, the Q factor was fudged to be no less than 150. This limit was never reached before the 25th mode of the strings; moreover, it will be seen later that the fractional delay filter used for the accurate tuning of the strings adds some damping in the high-frequency range, which compensates, to some extent, for the underestimation of damping in that range.

The next step is the detailed filter design. The method used here is similar to the one described in [7]: Matlab's *invfreqz* routine is used to design a filter based on the desired amplitude response. As with any other phase-sensitive filter design method, *invfreqz* gives its best result when designing a minimum-phase filter; for that reason, a minimum-phase version of the desired amplitude response is made first. This was achieved using the non-parametric method of folding the cepstrum to reflect non-minimum-phase zeros inside the unit circle [7]. The weight function for *invfreqz* is set to $1/f$, and the filter is designed with one zero and 300 poles by default. If the initial number of poles results in an unstable filter, the number is changed iteratively until a stable filter is achieved: this method led to stable filters for the first octave on the C_2 and D_3 cello strings. A filter with 300 poles may seem excessive, but a high-order filter proved necessary to ensure a good fit at the first few string modes, particularly for the C_2 string (this issue is further discussed in Section 3). Several attempts were made to design Finite Impulse Response (FIR), rather than Infinite Impulse Response (IIR), damping filters both by truncating the inverse FFT of the desired frequency response and by using Matlab's filter design toolbox. Both methods proved to be problematic, particularly for the shorter segment of the string, and the fit was never as good as the one obtained by *invfreqz*. It is

not claimed that one cannot design an equally suitable FIR filter for this application, simply that we failed to do so.

The designed damping filter was phase-equalised using Matlab's *iirgrpdelay* routine (a 16th-order filter was used here). The minimum-phase damping filter and the phase-equalising filter were then cascaded into an almost-linear-phase damping filter with the desired amplitude response. The phase-equalisation may not have been fully successful in making the filter linear-phase, but this turns out to be unimportant once the dispersion filter is added, since it involves much more significant phase shifts.

Finally, tuning was implemented using a combination of an integer-sample delay and an order-6 Farrow fractional delay [23] for each side of the string (totalling β/f_0 for the bridge side, and $(1-\beta)/f_0$ for the finger side). When a stiff string was to be modelled, tuning was postponed until after the design of the dispersion filter. In summary, the order of the filters for each segment of the string is as follows: damping filter, phase-equalising filter, dispersion filter (if a stiff string is being modelled), integer delay filter, and fractional delay filter.

Dispersion was accounted for using an all-pass filter, with a unit gain at all frequencies, which delays the signal in a frequency-dependent manner. The method used to design such a filter was based on a technique introduced by Abel and Smith [24], which makes a dispersion filter in the form of cascaded first-order all-pass filters. This method was later applied to the particular problem of a stiff string in [25].

In brief, in this method the frequency-dependent part of the group delay (total delay of a stiff string minus the linear-phase term corresponding to a pure delay) is broken down into segments of 2π area. Associated with each segment is a first-order all-pass filter with a pole placed at the centre of the corresponding frequency band. The pole radius sets the bandwidth of the group delay peak for each band, and in that way determines the trade-off between the smoothness of the final filter and its ability to track sudden changes in the desired group delay. The radius of each section is set so that within each band the minimum group delay (happening at the edges of the band) is equal to 0.85 times the maximum group delay (happening at the centre of the band). Ultimately the designed first-order sections are combined with their complex conjugates to produce real second-order all-pass filters. These second-order filters are cascaded and directly implemented into the loop filter without being converted to the transfer function form. The reason for this is to avoid round-off errors resulting in an unstable filter, a common problem for all-pass filters [26].

The original implementation proposed in [25] uses a first-order Newton's approximation to find the solution to the equation that gives the frequency of the poles (Equation (8) in [25]); but here the exact solution to that equation has been calculated. The first-order approximation gave a convincingly close approximation to the desired behaviour for the longer segment of the string (although, not surprisingly, never as good as the closed-form solution),

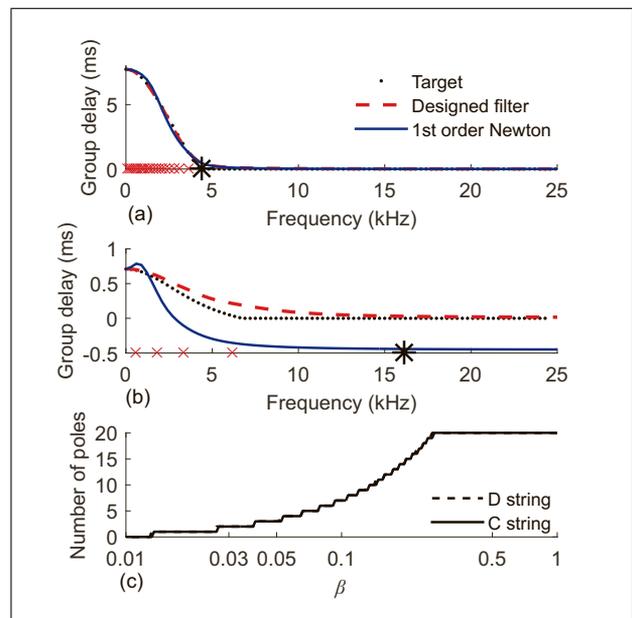


Figure 2. Group delay of the designed filter (dashed line) for the finger side (a), and the bridge side (b) of the open cello C_2 string compared to the desired response (dotted line), and a filter designed with first-order Newton's approximation (solid line). The crosses show the position of the poles used in the designed filter and the star shows the upper limit of the design frequency range. A constant group delay is assigned for the frequencies beyond that range. (c) shows the order of the dispersion filter for the C_2 and D_3 strings as a function of β .

but it proved to be problematic in designing the dispersion filter for the shorter segment of the string, at least for the way it was originally implemented in [25]. Figures 2a and 2b show the desired group delay behaviours against the results obtained from the exact solution and the first-order approximation, respectively for the short and long segments of the open C_2 string (β is here chosen to be 0.10).

Filters designed in this way give an almost constant group delay to all frequencies above the target frequency (marked by a star on the horizontal axis of Figures 2a and 2b), which results in a spike-like behaviour in the equivalent reflection functions (see Figure 3 and the following discussion). Time-domain details of this kind may be insignificant in producing audible effects as human ears are not too sensitive to phase, but they may affect the playability of a simulated bowed string by creating an unphysical disturbance at the bowing point. This can significantly compromise the accuracy of the model in predicting the playability of a bowed string. In this regard, a relatively high order (order-20) dispersion filter was often found to be necessary, especially for the finger side of the string. The order was reduced whenever an order-20 filter resulted in a design frequency range passing the Nyquist rate (common for the bridge side and for a small bow-bridge distance). The order of the dispersion filter for the C_2 and D_3 cello strings as a function of β is illustrated in Figure 2c — the two curves are so similar that they can hardly be dis-

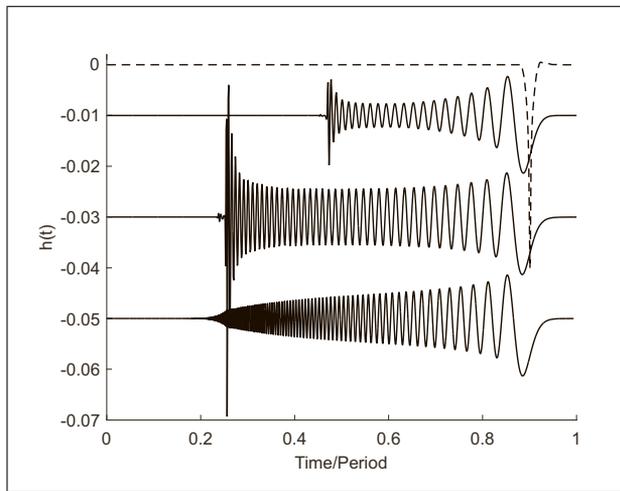


Figure 3. Equivalent reflection function (impulse response of the loop filter) designed for the finger side of a damped cello C_2 string, perfectly flexible (dashed line) and stiff (top solid line). The natural frequency of the string is 65.4 Hz, $\beta = 0.1$, frequency-dependent Q factor based on the data in Table I, bending stiffness is $6.2 \times 10^{-4} \text{ Nm}^2$, sampling frequency $6 \times 10^4 \text{ Hz}$. The middle solid line is the same as the top solid line except the Q factor of the string modes is assumed constant at 600, and the number of poles in the dispersion filter is increased from 20 to 40. The bottom solid line is the equivalent of the middle solid line but damping is modelled using the constant-Q reflection function of Equation (5), and dispersion is implemented based on the method proposed in [8]. Note the spike-like behaviour in the top solid line at non-dimensionalised time 0.47, and more vividly, in the middle solid line at non-dimensionalised time 0.25, resulting from frequencies above the design frequency of the dispersion filter.

tinguished in the plot. The dispersion filter was excluded whenever the filter order would become less than 2, which is the case for β smaller than 0.028.

The equivalent reflection function for the finger side of the open cello C_2 string is shown in Figure 3, both for a perfectly flexible and for a stiff string. Damping parameters for both plots are based on the data in Table I, and β is again set at 0.10. Even with an order-20 dispersion filter, some evidence of the spike-like behaviour can be seen at non-dimensional time 0.47 for the stiff string case. The plot also shows the result for a constant Q of 600 implemented using an order-40 filter. This may be compared with the bottom trace, which shows the corresponding result based on the earlier modelling (damping modelled using the constant-Q reflection function of Equation (5), and dispersion implemented based on the method proposed in [8]).

The inharmonicity of the n th partial of the full string is jointly defined by the inharmonicities for the two segments of the string. Having that in mind, for the cases where the bow/pluck is extremely close to the bridge the Nyquist rate may only cover the first few partials, leaving the higher partials of the full string with an effective inharmonicity that is less than the target value. As a practical fix for those cases, an inflated inharmonicity was given to the finger side of the string to compensate.

2.2. Coupling to the instrument body

The next stage of modelling is to couple the string to the body of the instrument. The vibrating string exerts a force on the bridge, which evokes a response from the body. That response will not in general be in the same direction as the applied force, so the body motion excites some motion of the string in the polarisation perpendicular to the original one. This makes it natural to treat the two effects together. The second polarisation of string motion can be treated by the method introduced in the previous subsection, with two additional travelling wave components and an identical set of reflection functions to describe the damping and dispersion. The body response at the bridge can be characterised in terms of a 2×2 matrix of frequency response functions, giving the components of body motion in the two planes in response to forces in those planes.

The frequency response function most commonly used is the admittance (or mobility): the velocity response to applied force. The matrix of admittances can be expressed in terms of the modal parameters of the body, by a standard formula. Define the direction X to be tangent to the bridge-crown for a violin or cello, and define the direction Y perpendicular to both the X -direction and the string axis. If $F_{X,Y}$ and $V_{X,Y}$ are the components of force and velocity in these two directions, then the admittance matrix is defined by

$$\begin{bmatrix} V_X \\ V_Y \end{bmatrix} = \begin{bmatrix} Y_{XX} & Y_{XY} \\ Y_{YX} & Y_{YY} \end{bmatrix} \begin{bmatrix} F_X \\ F_Y \end{bmatrix}, \quad (8)$$

where

$$\begin{bmatrix} Y_{XX} & Y_{XY} \\ Y_{YX} & Y_{YY} \end{bmatrix} = \sum_k \begin{bmatrix} \cos^2 \theta_k & \cos \theta_k \sin \theta_k \\ \cos \theta_k \sin \theta_k & \sin^2 \theta_k \end{bmatrix} \frac{i\omega u_k^2}{\omega_k^2 + i\omega\omega_k/Q_k - \omega^2}, \quad (9)$$

and where the k th mode has natural frequency ω_k , Q factor Q_k , mass-normalised modal amplitude at the string notch in the bridge u_k , and a “modal angle” θ_k defined as the angle of the principal direction of bridge motion in that mode with respect to the X -direction [1].

The first step to implement a realistic body model is to extract the relevant set of modal properties of an actual instrument. Calibrated measurements were carried out on the bass-side corner of the bridge on a mid-quality cello. A miniature hammer (PCB Model 086E80) and LDV (Polytec LDV-100) were used to measure the 2×2 admittance matrix. The strings were correctly tensioned, but during this measurement they were thoroughly damped (including their after-lengths) using small pieces of foam. Mode fitting was performed by an analysis method described in [27], using the Matlab function *invfreqs*. The method first involves modal extraction through pole-residue fitting, followed by an optimisation procedure allowing selection of the best sets of complex and real residues by minimising the mean of the modulus-squared deviation between measurement and reconstruction. This method was performed

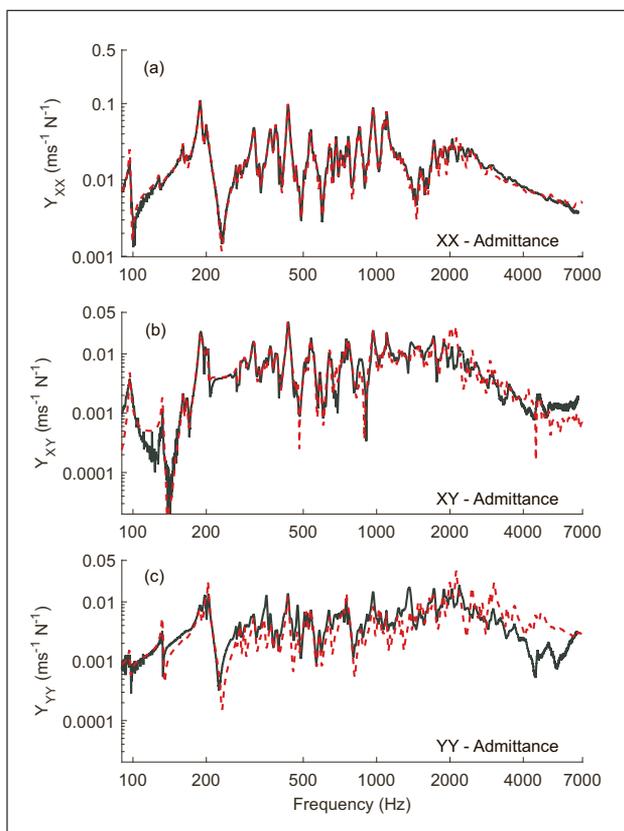


Figure 4. Measured admittances in the plane perpendicular to the string axis (green solid curve, colour online) and the fitted admittances to them (red dashed curve, colour online) for (a) XX admittance (b) XY admittance and (c) YY admittance. Note that the vertical scales are different in (a) and (b)-(c).

on Y_{XX} and Y_{YY} separately, and then modes that were recognisably the same for the two fittings were merged to give a final set of frequencies and Q factors. Modal masses and spatial angles were then optimised to give the best fit to all admittances.

To maintain the quality of fit the frequency range 0–90 Hz was included, but the modes falling within that range were later removed because these were all identified as fixture modes in which the cello moves essentially as a rigid body. Beyond 2 kHz, the modal overlap increases and the fitting process becomes increasingly unreliable. A statistical fit was then used, exactly as done earlier by Woodhouse [1] for the guitar. The procedure assigned 166 extra modes to the frequency range 2–7 kHz, using a random number generator to create modal frequencies with correct density and spacing statistics, as well as damping factors and modal masses with approximately correct statistical distributions. The resulting fit is compared to the measured admittances in Figure 4. The corresponding phase fits showed excellent fidelity up to 2 kHz although deviating a little at higher frequencies, especially for the XY admittance.

To implement the body dynamics in the model, each body mode is simulated as an independent resonator excited by the force exerted by the string at the bridge. It would be possible to include the body modes inside the

IIR loop filter of the bridge side, but it is useful to have direct access to the physical velocity of the bridge, so it was decided to implement them separately. This also gives a simple and efficient means to synthesise the radiated sound from the instrument. The complex amplitude of the k th mode at sample $i + 1$ can be calculated from its amplitude at sample i by

$$A_{k,i+1} = A_{k,i} e^{(i\omega_k - \omega_k/2Q_k)h} + hu_k^2 F_k, \quad (10)$$

where h is the time-step and F_k is the instantaneous force applied by the string by the incoming waves (in both transverse polarisations), projected in the principal direction of mode k ,

$$F_k = -2 \cdot Z_0 (v_{oX} \cos \theta_k + v_{oY} \sin \theta_k). \quad (11)$$

Here v_{oX} and v_{oY} are velocity waves sent from the excitation point towards the bridge in the X and Y polarisations β/f_0 seconds before the current time-step, and $Z_0 = \sqrt{T_0 m_s}$ is the characteristic impedance of the string.

The physical velocity of the bridge projected in the X and Y directions can be obtained by summing the contributions of all body modes,

$$\begin{aligned} V_X &= \Re e \left\{ \sum_k A_k \cos \theta_k \right\}, \\ V_Y &= \Re e \left\{ \sum_k A_k \sin \theta_k \right\}. \end{aligned} \quad (12)$$

These projected velocities then contribute to the history of v_{oX} and v_{oY} , after filtering by the bridge-side loop filter to give the actual velocity waves arriving back at the bowing/plucking point. For the finger side the incoming waves are calculated simply by filtering the history of the outgoing waves toward the finger by the finger-side loop filter. For cases when a single-polarisation simulation of the string was wanted, the terms in the Y -direction were omitted.

The schematic of the coupled string-body model for a single polarization of a plucked string is illustrated in Figure 5.

2.3. Additional details

On most stringed instruments, several strings are supported on a common bridge and are coupled to one another through that path. Although coupling happens between all such strings, the effect is much stronger if the tuning of the strings is close to unison or otherwise harmonically related. This effect has been known to instrument makers for a very long time, as is evident from the existence of sympathetic – but non-played – strings in many instruments such as the Norwegian Hardanger fiddle, the Indian Sarangi, or the Persian Rubab. Sympathetic strings can create a number of interesting musical effects, most famously the multi-stage decay arising from slight mistuning of pairs or triplets of nominal unison strings in the piano [28].

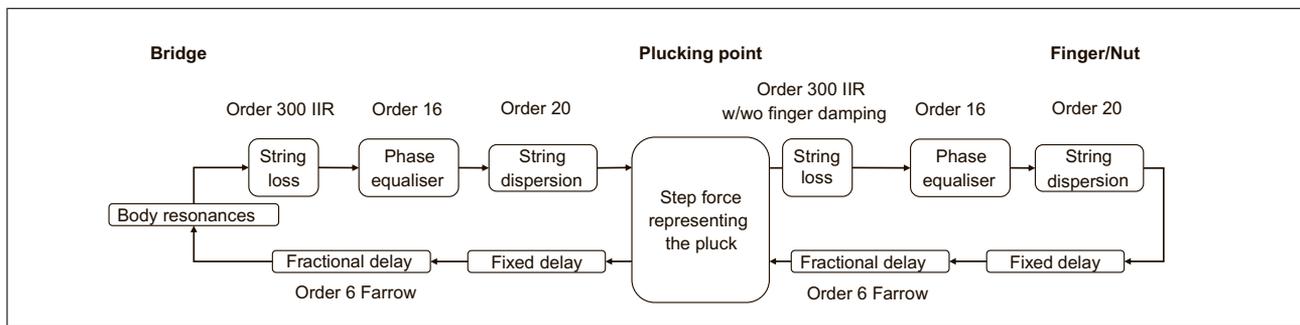


Figure 5. Schematic of the plucked-string model.

Such sympathetic strings can be straightforwardly included in the simulation model by adding the reaction force of all strings to Equation (11). Similar to the case for a single string, the contribution of the moving body adds to the reflected waves at the bridge, this time for all strings. Since the only excitation acting on the sympathetic strings is the moving bridge, they can be modelled with a single loop-filter describing the round trip wave propagation from the bridge to the finger and back.

For instruments like the cello, the strings pass over the bridge and join to the tailpiece. These after-lengths could be added to the model using the same method, except that they are terminated at a fairly flexible floating tailpiece rather than a rigid termination at the nut. Natural frequencies and mode shapes of a cello tailpiece can be found in [29], and they can be included in the modelling scheme exactly as the body modes were included. A computationally-cheaper alternative might be to measure the bridge admittance with the after-lengths undamped, and to include them implicitly into the model of the body. However, this would compromise the link between the model and the underlying physics and make it harder to explore the influence of, for example, changing a tailpiece mode frequency.

2.4. Simulating the pluck

The initial condition of an idealised plucked string is zero velocity, and non-zero displacement (and acceleration). In principle, it is possible to initialise the waveguides to produce arbitrary initial conditions; the values of the two travelling waves add to form the physical velocity at each point, so there are two degrees of freedom to set the desired initial velocity and acceleration [7]. Although that possibility was available, an alternative approach is used here.

An ideal pluck can be created by pulling a single point of the string sideways and then suddenly releasing it with no initial velocity: the force for such a pluck has a constant non-zero value F_P for $t < 0$, which suddenly drops to zero at $t = 0$. If this force is offset by an amount $-F_P$, the only effect is a fixed static offset in the displacement of the string, which does not matter in the context of linear theory since superposition can be used. (Note that this is quite a different effect from the *velocity* offset that would

arise if Equation (1) was not satisfied.) This allows a simple “trick” option for implementation: both travelling velocity waves can be initialised to zero values, and at $t = 0$ a constant force is applied at the plucking point which persists over the time of simulation. The direction of the step force corresponds to the angle of release of the pluck, and can be varied at will: this angle is used by guitar players to influence the tone color and the decay rate of the sound produced by the instrument (a comprehensive discussion of the topic can be found in [30]).

Such an ideal pluck is hard to achieve in reality: the closest one can get is by looping a thin wire around the string at the plucking point and gently pulling the wire until it breaks. Using a fingertip or a plectrum of finite size results in additional rounding of the shape of the string at the plucking point and hence in a low-pass filtering effect on the played note. The detailed interaction of a plectrum or fingertip with the string and the exact way the pluck is executed have a significant effect on the final sound of the instrument: this has been discussed in some detail in [31, 32].

3. Evaluating the accuracy of the plucked-string model

It is important to assess the accuracy of the simulation methodology described above. As a preliminary test the method was applied to guitar plucks, using the string and body properties from the earlier study by Woodhouse [1]. The results, not reproduced here, showed excellent agreement with the other synthesis methods explored in that study. The problems with the time-domain approach reported in that study are thus seen to stem from an insufficiently accurate implementation of the method, rather than from any fundamental shortcoming in the approach. This is reassuring, but it is not a test of the accuracy of the model: it merely compares different numerical approaches to solving the same model. What is needed is direct comparisons with measurement.

The techniques described above were applied to simulate 10 s of plucked sound for the first 12 notes on the C_2 and D_3 cello strings. The damping added by the finger of the player is included, except for the open strings. Some representative sound examples, for the simulated

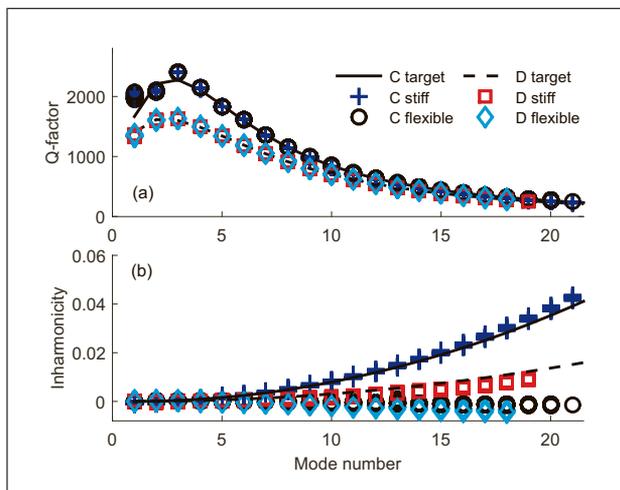


Figure 6. Trend of the Q factor (a) and inharmonicity (i.e. $[(f_n/nf_0)^2 - 1]$) (b) versus the string mode number for the stiff and flexible open C_2 and D_3 cello strings. All strings were terminated at rigid boundaries and the results are extracted from 10 s of simulated plucked response. β is varied in 20 steps and the results are overlaid.

open D_3 string, are available at [33], illustrating what happens when different features are progressively added to the model. Cases include a perfectly flexible string terminated at rigid ends, a stiff string terminated at rigid ends, a stiff string terminated at a realistic bridge and vibrating in a single polarisation, a stiff string terminated at a realistic bridge and vibrating in both polarisations, and finally the sympathetic strings are added. The response is the velocity wave on the string travelling towards the bridge, which is proportional to the transverse force applied by the string to the bridge. The signal that is converted to a sound file is a low-pass filtered version of that travelling wave, to simulate the radiation from the instrument's body, crudely, by treating the body as a pulsating sphere of roughly the right diameter (see Equation (6) of [11]).

The simulated results for the set of notes on the C_2 and D_3 cello strings were analysed to extract the frequency and Q factor of at least the first 15 string modes by the same method used earlier with experimental data. Figure 6 shows the extracted Q factors and inharmonicities (equal to Bn^2 in Equation (3) and calculated from $[(f_n/nf_0)^2 - 1]$ for each string mode) for the two open strings, with and without allowing for string stiffness. For the moment, an open string case with rigid end terminations is chosen to focus on the results of the damping and dispersion modelling. Figure 6 includes 20 different β values (i.e. different pluck-bridge distances). Ideally, both Q factor and inharmonicity should be independent of the plucking point, so that plots for different β values should overlay. This clearly is the case except for the first two string modes of the C_2 string, where slight variation can be seen. This variation vanishes almost entirely as soon as the bridge is turned from a rigid termination to a realistic flexible one.

The target trends for Q factor and inharmonicity from Figure 1 are also overlaid for both strings. Accurate track-

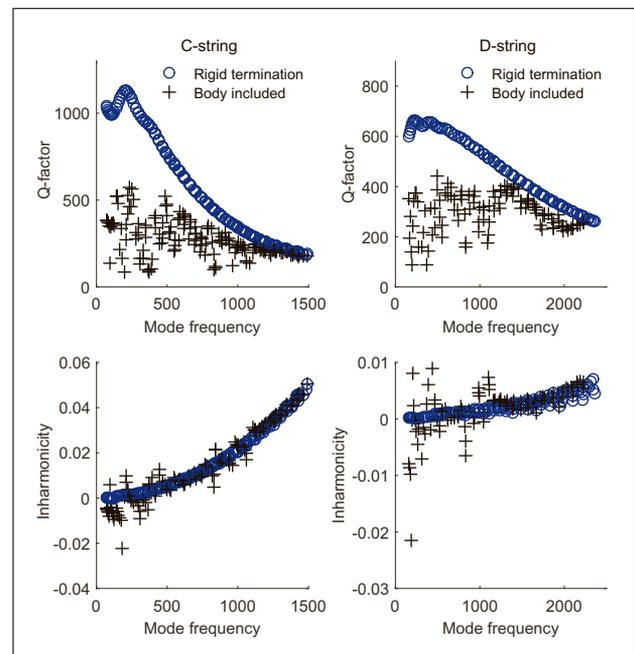


Figure 7. Trend of Q factor (upper plots) and inharmonicity (lower plots) versus the string mode frequencies for the stiff C_2 (left plots) and the stiff D_3 (right plots) cello strings. Circles show the case when the strings are terminated at rigid boundaries and plus signs show the case when the flexible body is included in the simulations. The first 11 semitones have been “played” on each string and the results were extracted from 10 seconds of a simulated plucked response.

ing of the desired Q factor is seen, but this could only be achieved by using a very high order damping filter; reducing the number of poles from 300 to 100 significantly degraded the final result. Inharmonicity in the “perfectly flexible” cases for both C_2 and D_3 strings shows some deviation from the expected zero value, caused by limitations of the phase-equalisation procedure, but the range of variation is almost negligible compared to the inharmonicity caused by stiffness. Note that the desired Q factor and inharmonicity trends are genuinely different for the C_2 and D_3 strings, so the plot for each stiff string should be only compared to its corresponding flexible one. It is satisfactory to see that the Q factors for both strings are not affected by the dispersion filter.

Figure 7 shows what happens to the simulated results when the body contribution is added to the model. Since it has already been demonstrated that the response of the string is not a function of the plucking point, the plots are only drawn for the smallest β value (equal to 0.02), to excite the largest number of string modes before the first missing harmonic appears (at $n \approx 1/\beta$); instead, the plot includes the first 11 finger-stopped semitones on each string. The equivalent results for the case of rigidly terminated strings are also included for comparison; string stiffness is included in both sets of simulations. The Q values are of course lower than those of the open strings, due to the additional damping from the finger. The Q factors and inharmonicities are both plotted against the string mode

frequency and are overlaid for different notes played on the same string.

As expected, once the body is included in the model the Q factors drop significantly and in a frequency-dependent manner. The frequencies of the string modes are perturbed compared to their counterparts obtained with rigid terminations, more severely at lower frequencies where veering is more likely to occur [34, 35]. The ceiling level of the Q factors for the modes of a string mounted on an actual cello does not quite reach the Q factor of the same string with rigid end terminations: for instance the highest Q factor among all partials for the C_2 string barely reaches 600, compared to 1200 achieved with rigid end terminations. The numbers are much lower than those in Figure 1 because finger damping has been added.

The next step is to compare the simulated coupled string-body model with its experimental counterpart. Figure 8 shows the simulated Q factors for the open C_2 and D_3 cello strings (terminated with rigid ends and with the body model) overlaid on experimental data obtained from the same cello whose bridge admittance was used to fit the modal properties. The results are in very good agreement with the numerical predictions, showing only very modest discrepancies. In any case, the exact values of the measured Q factors should not be over-interpreted: they will be sensitive to string excitation angle and exact tuning, as well as to the usual uncertainties in measuring vibration damping.

As another useful check for the simulation of string-body interaction, one can treat the model as an actual instrument with strings undamped and simulate the standard measurement of the bridge admittance by exciting the bridge with an impulse and measuring its velocity. Figure 9 shows the result of such assessment. Both polarisations of all four strings were included in the model, excited only via the bridge motion. The simulated bridge admittance in the X -direction is compared to the measured one, when all strings were free to vibrate. The plots are all to scale, and no modification has been made to match the two.

As one would expect, the general trend of the admittance for the strings-undamped case is similar to the strings-damped cases (earlier shown in Figure 4a), the only significant difference being sharp string resonances and antiresonances appearing in the strings-undamped version. Figure 9b is a zoomed version of a particular frequency range of Figure 9a: the “wolf note” area. The strongest body effect is around the wolf frequency, and it is interesting to see how the sympathetic strings interact with the body modes present in that frequency range. The 2nd harmonic of the G_2 string and the 3rd harmonic of the C_2 string both fall in that region. The two would coincide if the strings were perfectly flexible, but are slightly mistuned due to different inharmonicities. Both the experimental bridge admittance and the simulated one for the strings-undamped case are added to the plot, for comparison. It can be seen that the two strong modes falling on either side of the string resonances have been repelled by

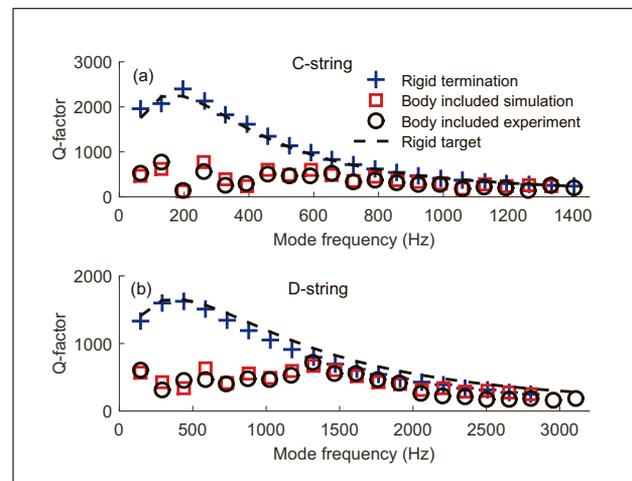


Figure 8. Q factor versus fundamental frequency for the open C_2 (a) and open D_3 (b) cello strings. Plus signs show the Q factor of the synthesised pluck with rigid terminations, squares show the same quantity when the coupling to the body modes are included and the circles show data measured on an actual instrument. The body modes were fitted to the bridge admittance of the same instrument.

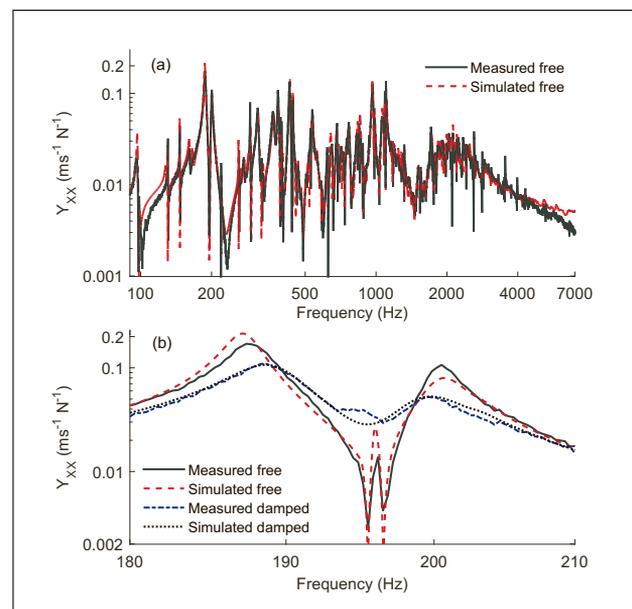


Figure 9. Simulated versus measured bridge admittance in the X direction when all four strings are free to vibrate (a), and a zoomed version of that plot covering only the “wolf-note” area (b). Both measured and simulated data for the strings-damped case are also included in the lower plot for comparison.

the reactive components of the string modes (see [34] for an explanation). These effects have been very well captured by the model.

Finally, Figure 10 shows the equivalent of Figure 6 but using the constant-Q reflection function of Equation (5) and the old implementation of dispersion proposed in [8]. This particular combination was used in many earlier studies, such as [17, 18]. Figure 10 shows the Q factor and inharmonicity of the open D_3 string, with and without dispersion and for 20 different β values. Note that the older implementation uses a constant-Q damping model (set to

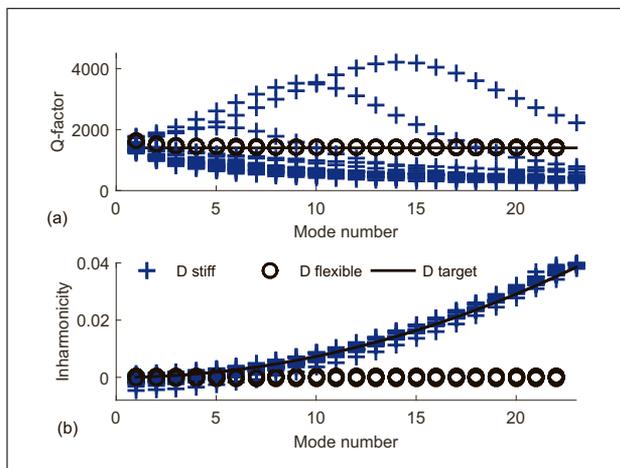


Figure 10. Trend of Q factor (a) and inharmonicity (i.e. $[(f_n/nf_0)^2 - 1]$) (b) versus the string mode number for stiff and flexible open D_3 cello string, based on the old implementation. The strings had rigid terminations and the results were extracted from 10 seconds of simulated plucked response. β was varied in 20 steps and the results are overlaid.

1800 here) and for that reason is not directly comparable to the results presented in Figure 6. The sampling rate to obtain the results of Figure 10 is set to 200 kHz (compared to 60 kHz used for this newer implementation), as used in some of the earlier studies.

It can be seen that the Q factor of a perfectly flexible D_3 string follows the intended constant value of 1800 fairly accurately. For the same simulation made on the C_2 string or with a lower sampling rate on the D_3 string (neither reproduced here), the Q factors of the first few string modes were slightly above the desired value. As was discussed earlier this effect is an artefact of how normalisation was carried out in the process of designing the filter. Gratifyingly, the inharmonicity of the perfectly flexible case stays very close to zero, more accurately than was the case for the newer implementation presented earlier.

Once the dispersion is included, the results are much less satisfying. Although the inharmonicity of the simulated plucks matches the desired trend very well, it drastically affects the Q factor of the partials, and it has also made the Q factor a sensitive function of β . Instability was also observed in some cases, which echoes earlier difficulties reported to synthesise a guitar pluck using this technique [1]. Including the body into the model alleviates the situation to an extent, but it is clear that the model presented here offers more flexibility and precision in tracking the target trend of damping.

4. Conclusions

A refined model of a plucked string based on time-domain simulation has been presented. Various details of the underlying physics have been incorporated into the model: the frequency-dependent damping of the string, an accurate implementation of dispersion, and the interaction of the string vibrating in two polarisations with a realistic bridge as well as the sympathetic strings supported on the

same bridge. Parameter values for the properties of the strings and body were extracted from measurements on a cello: the information about cello strings is itself a new contribution to the subject.

Using some sample results, it has been demonstrated that the model of the string precisely follows the target trend for the Q factors and dispersion. More importantly, the fully coupled model of the plucked string was compared to plucked notes of an actual instrument, which demonstrated the ability of the model to produce a response with very similar Q factors to the experiments. The simulated bridge admittance when all strings were either damped or free to vibrate was also compared to measurements. The results were almost indistinguishable for the strings-damped case. Finally, it was verified that the effect of sympathetic strings and their interaction with the body modes is very well captured by the model.

These results demonstrate that wave-based models can indeed simulate plucked strings with comparable fidelity to modal-based methods (see for example [35, 36]). This may seem a rather minor contribution, since the modal methods are already available. If the only purpose were to simulate plucked strings, this would be a fair objection. However, the model has been developed in a form suitable for extension to the case of bowed excitation of the strings, and the details of that case are explored in a companion paper [2]. For bowed strings, the relation to the modal approach reverses: while it is indeed possible to study bowed strings by a modal method (see for example [37]), the non-linear nature of the friction force makes a time-domain approach more natural and intuitive. As friction models become more sophisticated in the search for physical accuracy, this distinction is likely to become stronger, and it is hoped that the model presented here will form a strong foundation for such studies.

References

- [1] J. Woodhouse: On the synthesis of guitar plucks. *Acta Acustica united with Acustica* **90** (2004) 928–944.
- [2] H. Mansour, J. Woodhouse, G. Scavone: Enhanced wave-based modelling of musical strings. Part 2: Bowed strings. *Acta Acustica united with Acustica* **102** (2016) 1094–1107.
- [3] O. Darrigol: The acoustic origins of harmonic analysis. *Archive for history of exact sciences* **61** (2007) 343–424.
- [4] J. D’Alembert: Recherches sur la courbe que forme une corde tendue mise en vibrations. *Histoire de l’Académie Royale des Sciences et Belles Lettres* **3** (1747) 214–249.
- [5] M. E. McIntyre, J. Woodhouse: On the fundamentals of bowed string dynamics. *Acustica* **43** (1979) 93–108.
- [6] M. E. McIntyre, R. T. Schumacher, J. Woodhouse: On the oscillations of musical instruments. *The Journal of the Acoustical Society of America* **74** (1983) 1325–1345.
- [7] J. O. Smith: *Physical audio signal processing: For virtual musical instruments and audio effects*. W3K Publishing, 2010.
- [8] J. Woodhouse: On the playability of violins. Part I: Reflection functions. *Acustica* **78** (1993) 125–136.

- [9] R. T. Schumacher, J. Woodhouse: The transient behaviour of models of bowed-string motion. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **5** (1995) 509–523.
- [10] H. Jarvelainen, M. Karjalainen: Perceptibility of inharmonicity in the acoustic guitar. *Acta Acustica United with Acustica* **92** (2006) 842–847.
- [11] J. Woodhouse, E. K. Y. Manuel, L. A. Smith, A. J. C. Wheble, C. Fritz: Perceptual thresholds for acoustical guitar models. *Acta Acustica united with Acustica* **98** (2012) 475–486.
- [12] H. Fletcher, E. D. Blackham, R. Stratton: Quality of piano tones. *The Journal of the Acoustical Society of America* **34** (1962) 749–761.
- [13] M. Campbell, C. Greated: *The musician’s guide to acoustics*. 2nd ed. Oxford University Press, New York, 1994.
- [14] E. Ducasse: On waveguide modeling of stiff piano strings. *The Journal of the Acoustical Society of America* **118** (2005) 1776–1781.
- [15] R. Pitteroff, J. Woodhouse: Mechanics of the contact area between a violin bow and a string. Part I: Reflection and transmission behaviour. *Acta Acustica united with Acustica* **84** (1998) 543–562.
- [16] J. Woodhouse, A. R. Loach: Torsional behaviour of cello strings. *Acta Acustica united with Acustica* **85** (1999) 734–740.
- [17] P. M. Galluzzo: On the playability of stringed instruments. Dissertation. Engineering Department, University of Cambridge, Cambridge, UK, 2003.
- [18] R. Pitteroff, J. Woodhouse: Mechanics of the contact area between a violin bow and a string. Part II: Simulating the bowed string. *Acta Acustica united with Acustica* **84** (1998) 744–757.
- [19] F. C. Tao: personal communication. 2015.
- [20] C. Valette: The mechanics of vibrating strings. – In: *Mechanics of musical instruments*. A. Hirschberg (ed.). Springer-Verlag, Vienna, 1995, 115–183.
- [21] J. Woodhouse: Plucked guitar transients: Comparison of measurements and synthesis. *Acta Acustica united with Acustica* **90** (2004) 945–965.
- [22] S. M. Saw: Influence of finger forces on string vibration. Diploma Thesis. Engineering Department, University of Cambridge, Cambridge, UK, 2010.
- [23] C. W. Farrow: A continuously variable digital delay element. *Circuits and Systems*, 1988., IEEE International Symposium on, 1988, IEEE, 2641–2645.
- [24] J. S. Abel, J. O. Smith: Robust design of very high-order allpass dispersion filters. *Proceedings of the International Conference on Digital Audio Effects*, Montreal, Quebec, Canada, 2006, 13–18.
- [25] J. S. Abel, V. Valimaki, J. O. Smith: Robust, efficient design of allpass filters for dispersive string sound synthesis. *Signal Processing Letters, IEEE* **17** (2010) 406–409.
- [26] J. O. Smith: *Introduction to digital filters: with audio applications*. Vol. 2. W3K Publishing, 2007.
- [27] J. Woodhouse, R. S. Langley: Interpreting the input admittance of violins and guitars. *Acta Acustica united with Acustica* **98** (2012) 611–628.
- [28] G. Weinreich: Coupled piano strings. *The Journal of the Acoustical Society of America* **62** (1977) 1474–1484.
- [29] E. Fouilhe, G. Goli, A. Houssay, G. Stoppani: Vibration modes of the cello tailpiece. *Archives of acoustics* **36** (2011) 713–726.
- [30] B. Scherrer: Physically-informed indirect acquisition of instrumental gestures on the classical guitar: Extracting the angle of release. Dissertation. Dept. of Music Research, McGill University, 2013.
- [31] D. Chadeaux, J. L. Le Carrou, B. Fabre: A model of harp plucking. *The Journal of the Acoustical Society of America* **133** (2013) 2444–2455.
- [32] G. Cuzzucoli, V. Lombardo: A physical model of the classical guitar, including the player’s touch. *Computer Music Journal* **23** (1999) 52–69.
- [33] Synthesized sound samples (cello). http://mt.music.mcgill.ca/hosseinn/research!/mus_acc/thesis_mat/. Accessed: 2016-03-11.
- [34] C. E. Gough: The theory of string resonances on musical instruments. *Acustica* **49** (1981) 124–141.
- [35] J. Woodhouse: On the synthesis of guitar plucks. *Acta Acustica united with Acustica* **90** (2004) 928–944.
- [36] V. Debut, J. Antunes, M. Marques, M. Carvalho: Physics-based modeling techniques of a twelve-string portuguese guitar: A non-linear time-domain computational approach for the multiple-strings/bridge/soundboard coupled dynamics. *Applied Acoustics* **108** (2016) 3–18.
- [37] O. Inacio, J. Antunes, M. C. M. Wright: Computational modelling of string-body interaction for the violin family and simulation of wolf notes. *Journal of Sound and Vibration* **310** (2008) 260–286.