

Audio-Rate Modulation of Physical Model Parameters

Edgar Berdahl

226 School of Music

Louisiana State University

Baton Rouge, LA, 70803, USA

edgarberdahl@lsu.edu

ABSTRACT

Audio-rate modulation of the parameters of physical models is investigated. For example, the resonance frequency of a simple resonator can be modulated smoothly at an audio rate to produce a brighter tone. As with traditional frequency modulation (FM) signal models, complex spectra can be produced through variation of the resonator's frequency.

It is possible to create such hybrid physical/signal models in which the physical portion conserves energy despite the modulation of its parameters. However, most interesting hybrid models will have some nonpassive characteristics. Example models are developed, whose vibrations can be felt and interacted with using a haptic force-feedback device. This technique can make exotic systems tangible that would not normally be found in nature. Hence, if one considers the process of designing a physical model a kind of "dematerialization" process, then one could consider that the present technique is physically "rematerializing" novel audio-rate-modulated virtual objects.

1. INTRODUCTION

Physical models have been employed for decades to synthesize sound with computers. The most basic physical modeling approach is to study the physics of a musical instrument, and then simulate the physical equations in a computer to synthesize sound [1, 2, 3, 4]. However, besides merely imitating prior musical instruments, new virtual instruments can be designed with a computer by simulating the acoustics of hypothetical situations [5], creating a "metaphorisation of real instruments." Of particular importance is also that sounds generated using physical models tend to be physically plausible, enhancing the listener's percept due to familiarity [6, 7].

That same ethos can be applied to the creation of digital musical instruments even with techniques besides physical modeling [8]. In other words, the idea is that by targeting sounds that share psychoacoustical characteristics of acoustic musical instruments, sound designers and composers will be able to "create powerful, expressive and evocative sounds" with signal-based sound synthesis [8].

1.1 Discussion of Continuous-Time Models

In order to make this work as widely applicable and general as possible, the models are presented in the continuous-time domain. These models describe what a physical system would do if its parameters were varied continuously. *If a model is not passive in the continuous-time domain, then it will most likely still not be passive when discretized*, and that is what is studied first and foremostly in this work. It is left to the reader as an exercise to decide which method to employ for discretizing the continuous-time differential equations to apply this approach in the future.¹

1.2 Control-Rate Modulation of Parameters

Physical modeling parameters can be adjusted continuously at control rates to simulate interactions with human players. For example, a sensible input to a physical model could be the position of a bow, and it would be natural to modulate or adjust this input parameter at a control rate to synthesize the sound of a bowed string [9]. Many researchers have also designed physical models for which the parameters were adjusted continuously, but nevertheless at a control rate [10, 3]. This kind of control-rate modulation makes the physical model emulate a time-varying system that is not guaranteed to be stable. However, in practice with many physical models, slow modulations tend not to drive the models unstable.

1.3 Audio-Rate Modulation of Parameters

In contrast, the concept of *audio-rate* modulation means that the physical model parameters are adjusted at rates faster than control or gestural rates. Examples include adjusting the fundamental frequency of a string according to a pseudorandom input audio signal or sinusoidally adjusting the resonance frequency of a resonator at 300Hz. From a compositional perspective, these examples are very intriguing because these sounds have not historically been heard by the wider public: it's very challenging to build a real string whose fundamental frequency is varied pseudorandomly, and it's hard to conceive how one might build a mechanical resonator whose frequency could be significantly varied at 300Hz without side effects. Thus, the concept of audio-rate modulation is intriguing. However, since

¹ The authors have discretized the models in order to simulate their dynamics using a computer; however, the only instabilities observed have been at low frequencies, suggesting that these instabilities are in fact due to the continuous-time model being unstable and not as a result of the nature of the discretization, and it is precisely those instabilities that are the subject of the present paper.

the modulations generally happen at fast speeds, audio-rate modulations are less likely to be stable (except for a specific class of passive models identified below in Sections 2.1-2.2). Therefore, when making music using audio-rate modulation of physical model parameters, care needs to generally be taken to avoid choosing unstable settings.

1.4 Tangibility

A further advantage of this approach is that it can make signal-based algorithms tangible with a force-feedback haptic device, since physical models are arguably the best way to program haptic devices [4]. The physical model allows a way for a user to interact with a signal-based algorithm as rendered by the physical model.

Alexandros Kontogeorgakopoulos et al. have completed related prior pioneering work in making signal-based algorithms tangible by way of *haptic audio effects* [11, 12]. In this work, a musical audio or speech signal is fed into a physical model as an input to be processed by the physical model. Kontogeorgakopoulos has derived ways of employing physical models to produce similar psychoacoustical effects as certain signal-based audio effects algorithms, thereby making the effects tangible [13]. The present work differs in that the application is not signal processing of an input musical audio or speech signal but rather direct sound synthesis.

2. PASSIVE MODELS

2.1 Passivity

If physical model parameters are modulated in a passive way, then energy cannot be created by the model or its modulation. Consequently, subject to assumptions about the haptic hardware and software implementation, the net system will be stable [4]. Although stability is not necessarily a requirement for musical applications, it can be convenient if a system does allow for implementations that are guaranteed stable. For this reason, we introduce continuous-time models that are guaranteed to be stable in the following section.

2.2 Modulation of Losses

As a mechanical musical instrument is struck, plucked, bowed, etc., energy is imbued into it. It will continue to vibrate until *losses* remove this energy, eventually allowing the instrument to essentially come to rest. There are many sources of loss, including friction and damping, that cause the energy of vibration to be converted into heat, etc. When such losses are incorporated into a continuous-time physical model, then any terms providing such pure losses can be modulated in real time without destabilizing the model, as long as the model does not incorporate any external sources of energy and as long as the loss parameters do not change sign. The reason for this is that pure losses instantaneously dissipate energy; therefore, increasing or decreasing the (positive) loss coefficients should not be able to destabilize a (passive physical) system that is initialized with only a finite amount of energy [14].

Consider a network of interconnected masses and springs, for which the springs incorporate some linear damping loss. If the coefficient describing this damping is suddenly reduced to zero, resulting in an ideal spring, then this cannot destabilize the system since it does not contain any external sources of energy and is initialized with only a finite amount of energy. Conversely, if the coefficient describing this damping is suddenly increased, then it, because it implements a loss, will either decrease the time before the system eventually comes to rest, or under some circumstances it will cause no change in the energy decay time (e.g. if this energy is not accessible to the loss).

2.2.1 Example: Modulation of Damping in a 1DOF Mechanical Resonator

A mechanical resonator can be formed by connecting a mass m to ground via a spring (k, R) as depicted in Figure 1. If the mass is displaced, then it will tend to oscillate with respect to ground. To facilitate connection to a haptic device, the model also incorporates a plectrum, which provides for haptic interaction with the resonator (see Figure 1).




Figure 1. A resonator is formed by a mass m connected to ground via a spring (k, R). A human can pluck this resonator by way of a virtual plectrum and haptic device.

This approach was first tested using a single resonator, but the results were extended to a more complex model in order to obtain a more interesting sound. The model was a perhaps oversimplified, but nonetheless intriguing, modal model of four gamelan bars. The model did not incorporate beating or any significant amplitude modulation. The damping loss $R(t)$ for the gamelan bars was then configured to turn on and off at frequencies close to 4 Hz. In other words, the damping parameter $R(t)$ was switched back and forth between $0.3 \text{ N}/(\text{m}/\text{s})$ and almost no damping (i.e. close to $0 \text{ N}/(\text{m}/\text{s})$), as demonstrated in the example sounds files:

- *GamelanPassiveDampMod.wav* and
- *GamelanPassiveDampMod2.wav*.

The result was a gamelan-like sound and elementary haptic interaction, for which the sound evidenced noticeable amplitude modulation at frequencies close to 4 Hz. The waveform for one strike on one of the gamelan bars is shown in Figure 2.




Figure 2. Displacement of a virtual gamelan bar when plucked via a haptic device and subjected to passive damping modulation.

Further passive models can similarly be created by simply modulating pure losses such as damping in other models. However, some limitations in this technique can be observed. Its effect on the sound is limited by its passivity—as soon as the increased damping affects the sound, it also removes energy from the physical model, thereby bringing the model closer to rest. Therefore, in general when long decay times are desired, passive damping modulation is *not* an effective approach for employing external signals to modulate the sound of physical models, unless an additional energy source is provided that can sustain the virtual vibrations.

2.3 Nonpassive Damping Modulation

The approach from the previous section can be extended by allowing the modulating signal to make damping parameters become negative as well as positive. A “negative damper” is dangerous as it creates energy locally in a model and can destabilize the model; however, in some situations, it can be balanced by a sufficient amount of positive damping at a later time. In making the example sound *GamelanNonpassiveStrongRM.wav*, the damping parameter $R(t)$ was switched back and forth between -5.64 N/(m/s) and 6.36 N/(m/s) at 60 Hz. Therefore, on average the damping parameter was positive. The result was a somewhat Gamelan-like tone with a significantly long decay time but featuring a rapid amplitude modulation at 60 Hz. The sound could be described as a kind of ring modulation of the Gamelan sound. Such a model with large negative damping values could go unstable depending on the precise contents of the model; however, we found the model to be stable when connected to a plectrum and haptic device as depicted in Figure 1. The reader is encouraged to try out the same demonstration patch *GamelanDampingMod~.maxpat* to explore the dynamic behaviors discussed in Section 2.

3. NONPASSIVE MODELS

Most approaches for modulating the parameters of a physical model will result in nonpassive behavior, which may or may not be problematic for a specific application. In this section, some approaches are described using example models that provide useful results.

3.1 Audio-Rate Modulation of the Resonance Frequency of a Mechanical Resonator

While the elemental haptic interaction with the resonator shown in Figure 1 is intriguing, the sound is perhaps a bit bare. It sounds like a transient followed by an exponentially decaying sinusoid. Therefore, it is interesting to consider modulating the parameters of the resonator. In Section 2, modulating the damping was already considered; therefore, here we consider modulating the resonance frequency $f_0(t)$ of the resonator instead. One way to implement this is to allow the stiffness $k(t)$ and mass $m(t)$ of the resonator to be modulated, as indicated by the thin red arrows in Figure 3, while the damping parameter R is held constant. For this work, the resonator is implemented using a phasor approach with equivalent mass $m(t)$ as described further in Appendix A [15, 16].




Figure 3. Modulating the stiffness $k(t)$ and mass $m(t)$ of a mechanical resonator model to change its frequency.

3.2 Sinusoidal Frequency Modulation

To make the timbre more rich, the resonance frequency $f_0(t)$ can be varied sinusoidally [17].² The “Simple FM” synthesis program suggests employing the following equation to determine the frequency $f_0(t)$ of the resonator [3]:

$$f_0(t) = f_c \cdot \left(1 + d \sin(2\pi f_m t)\right), \quad (1)$$

where f_c is the carrier frequency, f_m is the modulation frequency, and d is the frequency deviation d .

The reader can hear the effect of the sinusoidal frequency modulation by comparing the sound files *resonator.wav* with *resonator-SFM.wav*. The frequency of the unmodulated resonator (i.e. f_c) is tuned to MIDI note 52 ($E3$ at 164.8Hz), and the magnitude response for a window near the beginning of the plucked resonator sound is shown in thick blue in Figure 4. For the frequency-modulated waveform *resonator-SFM.wav*, the tone is much brighter due to $d = 0.95$ and $f_m = 4.03f_c$. Because f_m/f_c is not tuned precisely to an integer, some beating is audible; however, because the modulation frequency is tuned almost harmonically, partials nearly corresponding to an (odd) harmonic

² Alexandros Kontogeorgakopoulos has explored a similar technique. Instead of implementing a model that gives precisely the same output as the FM synthesis algorithm, he has designed a series of physical models, whose fundamental frequencies vary over time due to various simulated physical influences [18].

series can be clearly observed in Figure 4 in thin red. Consequently, the timbre is brighter with sinusoidal frequency modulation.




Figure 4. Magnitude response of a simple plucked resonator (e.g. $d = 0$ in thick blue) overlaid over a frequency-modulated plucked resonator (e.g. $d = 0.95$ and $f_m = 4.03f_c$ in thin red). Note that the energy near 0Hz comes as a result of the model’s excitation via a virtual plectrum.

Although this model was not always passive, it is still musically useful. When the oscillator is *not* in contact with the plectrum (see Figure 4), the model is passive and sounds like an exponentially decaying FM synthesizer. In contrast, when the oscillator *is* in contact with the plectrum, the energy in the oscillator tends to slowly grow gradually, at a rate proportional to how hard the user is pressing into the oscillator (see Figure 10). This interaction is intriguing and musically useful.

3.3 Random Frequency Modulation

In contrast with sinusoidal modulation, the frequency of the resonator can alternatively be *randomly* modulated. The `rand~` object in Max is useful for this purpose. `rand~` outputs a noise signal that is bandlimited at frequency f_B . It does so by linearly interpolating its output between -1 and 1. Consequently, by introducing `rand~` using the same scaling constant d , it becomes possible to modulate the frequency of the resonator randomly. If it is desired to prevent the resonance frequency from becoming smaller than some f_{min} , then the following equation can be employed to determine $f_0(t)$:

$$f_0(t) = \max \left(\left| f_c \cdot \left(1 + d \cdot \text{rand}^{\sim}(f_B) \right) \right|, f_{min} \right), \quad (2)$$

for which f_{min} may be set to a value such as 20Hz to prevent subsonic sound.

If f_B is set to a low bandwidth such as $f_B = 5\text{Hz}$, and if $d \approx 0.1$ or similar, the resonance frequency will sound as if it is slowly wandering. This has an intriguing sonic character and can be performed expressively using a haptic device. The changes in resonance frequency can be felt, particularly when the resonance frequency is either becoming relatively low or is ceasing to be relatively low. Conversely, if f_B is set much higher, such as on the order of 500Hz, then the resonator will assume a more noise-like character. The model `pluck_a_resonator~.maxpat` included

with this paper can be used to experiment with the techniques from both this section and Section 3.2.

3.4 Plucked Strings

The same approaches outlined for resonators can also be applied to vibrating strings. Consider the simple plucked string model depicted in Figure 5. It can be adapted for real-time control of the fundamental frequency by adding a virtual slide bar as shown in Figure 6, for which the string is assumed to be arbitrarily long yet still tensioned appropriately. The position of the slide bar determines the point




Figure 5. Simple plucked string model.

at which the leftward-going wave is reflected back toward the right, hence setting the fundamental frequency of the portion of string that the listener hears (see Figure 6). The Max patch `pluck_a_string~.maxpat` included with this paper allows one to experiment with sinusoidal, sawtooth, and random frequency modulation of vibrating strings.




Figure 6. Arbitrarily long plucked string with slide bar for changing pitch.

3.4.1 Sinusoidal Frequency Modulation of Strings

Sinusoidal frequency modulation tends to make the string sound brighter, similarly to the case in Section 3.2.³

3.4.2 Sawtooth Frequency Modulation of Strings

Sawtooth frequency modulation with sufficiently large f_m is similar but has a more grainy character due to the transients in the sawtooth signal. In contrast, at such slow modulation frequencies (e.g. $f_m < 2\text{Hz}$ and similar), the sound is quite different. The instruments do not sound so much like strings anymore. Their pitch variations are almost dizzying, and the transient at the end of each sawtooth cycle has a perceptually similar outcome to re-energizing the strings over and over. Illustrative examples can be heard in the file `SawtoothStrings.mov`, which further demonstrates the advantage of locking the phases of multiple sawtooth waveform generators together in the case of multiple strings simultaneously.

³ For this model, the fundamental frequency of the string cannot be set negative so an absolute value function (or similar) must be incorporated into (1).

3.4.3 Random Frequency Modulation of Strings

The effect of random fundamental frequency modulation on the initial response of a plucked virtual string with $f_c = 36.7\text{Hz}$ (MIDI note 26, i.e. D_2) is represented in Figure 7. The unmodulated case with $d = 0$ is shown in thick blue, and the modulated case with $d = 0.15$ and $f_B = 6000\text{Hz}$ is shown in thin red (see Figure 7). Because f_B is set so large, the spectrum of the randomly modulated string is quite broadband, although clearly some of the lowest harmonics are still distinguishable above the noise-like sound (see the lowest seven or eight harmonics of the thin red spectrum in Figure 7). Readers may wish to listen to closely related sound examples in the included file *string_ex_MIDI_25.wav* or the video *RandomStrings.mov*.

3.4.4 Rhythm Generation using very low Fundamental Frequencies

One interesting usage of strings tuned to very low fundamental frequencies (e.g. 0.4Hz, 1Hz, etc.) is that they can be employed to generate rhythms. However, such rhythms can be performed using higher pitches by first initializing the strings at typical audio pitches. For example, if the strings are tuned to a chord, plucked, and then immediately dropped to very low fundamental frequencies, the audio content in the strings will slowly bounce back and forth creating rhythms with the tonal content originally stored in the strings. This can be observed in the sounds *subsonic-freq-4strings.wav* and *subsonic-freq-1string.wav*, the first part of which is shown in Figure 8. This example demonstrates how physical modeling can be employed to synthesize sounds that would be very difficult to create in the real world even using (modified) acoustic musical instruments in a scientific laboratory.




Figure 8. String that was plucked at $t=0$ sec when tuned to a note with a typical “audible” fundamental frequency, which was then dropped to about 0.4Hz starting at time $t=1.7$ sec.

4. MUSIC COMPOSITION *TRANSMOGRIFIED STRINGS*

The above methods have been applied to vibrating strings in the creation of the music composition *Transmogrified Strings*. For the composition, eight force-feedback sliders are controlled from a laptop to provide interaction with virtual vibrating strings and to produce eight output channels of audio. The haptic device housing the eight force-feedback sliders was custom made for performing the composition. The device consists of a clear, plastic acrylic box that houses four FireFader boards [19] (see Figure 9).

To emphasize the physical nature of the interaction, a video camera showing the view from Figure 9 is projected onto a large screen at the venue. The video delay should be as short as possible to promote coherence between the sound and the video.




Figure 9. Bank of four FireFader devices, including eight motorized faders total.

4.1 Philosophy

Haptic force-feedback devices can provide an exquisite tangible connection to software, but many questions still remain about how to program these devices. In music, another way of framing these questions is to ask, “How can haptic force-feedback devices best enable and support expressive new music performance?” *Transmogrified Strings* aims to help answer this question by investigating it in the context of plucked string sounds. For the composition, a series of virtual plucked string instruments have been designed, which have been made tangible (e.g. “rematerialized”) via the FireFader haptic device.

4.2 Programme Notes

Oxford Dictionaries defines *transmogrify* as “transform, especially in a surprising or magical manner,” which is the basic concept of *Transmogrified Strings*. It aims to surprise the listener with sounds that are both new yet uncannily familiar. In each section, virtual plucked string instruments are transformed via a specific kind of operation. For example, strings can be tuned as low as 0.5Hz or as high as the upper bounds of human hearing. The virtual strings retain their tangible character even as the sound changes drastically, and the feel of the instruments changes too, which in turn affects the performer’s rematerialized interaction with the virtual strings. The composition is organized into various sections, each of which is preceded and punctuated by the strumming of a harp. As the strings are *transmogrified* differently in each section, they are specifically

- I. *solemnified* (typical plucked strings playing a solemn melody),
- II. *demystified* (the strings are enlightened),
- III. *vivified* (the strings are plucked while gradually increasing their pitches and then instantly reducing their pitches again),




Figure 7. Magnitude response of a plucked string (e.g. $d = 0$ in thick blue) overlaid on top of a randomly frequency-modulated plucked string (e.g. $d = 0.15$ in thin red).

- IV. *solidified* (the string masses greatly increased, causing their pitches to become subsonically low to create rhythms), and finally
- V. *declassified* (a string is made to fall apart into individual, disconnected masses).

5. CONCLUSION

In this foray into the audio-rate modulation of physical models, it was decided to focus more on producing intriguing sounds and physical interactions rather than producing models which would be stable under all conditions. The reader is encouraged to download the audio samples, video samples, and Max patches associated with this paper for further evaluation of the models from a personal perspective.⁴

Several models for audio-rate modulation of the parameters of physical models were implemented and tested. In general, it was observed that slower audio-rate modulation, such as modulation with bandwidths less than 20Hz, were more likely to be stable than audio-rate modulation at faster rates. However, damping modulation was shown to be capable of implementing consistently passive models despite the audio-rate modulation, although the application of such models was limited to tones with shorter decay times.

Further models for audio-rate modulation, such as frequency modulation of physical models, were tested that were not guaranteed to be stable because they were non-passive. However, they were still useful for many parameterizations. This includes frequency modulation of a resonator. The implementation discussed in the appendices was determined to be applicable over a wide range of parameters.

Subjectively, *sinusoidal* frequency modulation was useful for brightening the timbre of low-order models at low computational expense and creating new sounds. *Sawtooth* and *random* frequency modulation were also subjectively intriguing, paving the way for creating more new sounds of unusual character. Such methods were employed in creating the music composition *Transmogrified Strings* for eight haptic devices.

5.1 Final Words on Tangibility

Overall the techniques described in this paper can make exotic systems tangible that would not normally be found in nature. Hence, if one considers the process of designing a physical model a “dematerializing” process, then one could consider that the present technique is “rematerializing” new kinds of audio-rate-modulated virtual objects. For example, a haptic interaction with an FM-modulated resonator (ref. Section 3.2) is depicted in Figure 10. While the plectrum is in contact (i.e. when the force in Figure 10 (bottom) is non-zero), the user can not only feel the vibrations of the resonator, but these vibrations are even strong enough to significantly modulate the position of the haptic device itself (see Figure 10, above in red) enabling a rich kind of bidirectional musical interaction—the user not only controls the physical model to produce intriguing sounds, but the model also exerts some control over the user. This interaction is the motivation for the development of the models in this paper.




Figure 10. Above: position of a frequency-modulated resonator (in blue) and a haptic device (in red), below: force between the frequency-modulated resonator and the haptic device as exerted by the plectrum.

6. REFERENCES

- [1] C. Cadoz, A. Luciani, and J.-L. Florens, “Synthèse musicale par simulation des mécanismes instrumentaux,” *Revue d’acoustique*, vol. 59, pp. 279–292, 1981.

⁴ <https://ccrma.stanford.edu/~eberdahl/ARMPMPExamples.zip>

- [2] J. Smith III, “Synthesis of bowed strings,” in *Proceedings of the International Computer Music Conference*, Venice, Italy, 1982.
- [3] C. Roads, *The Computer Music Tutorial*. Cambridge, MA: MIT Press, February 1996.
- [4] E. Berdahl, J.-L. Florens, and C. Cadoz, “Using physical models is necessary to guarantee stable analog haptic feedback for any user and haptic device,” in *Proceedings of the 8th Sound and Music Computing Conference*, Padua, Italy, July 6-9 2011.
- [5] N. Castagne and C. Cadoz, “Creating music by means of ‘physical thinking’: The musician oriented Genesis environment,” in *Proc. 5th Internat’l Conference on Digital Audio Effects*, Hamburg, Germany, Sept. 2002, pp. 169–174.
- [6] N. Castagné and C. Cadoz, “A goals-based review of physical modeling,” in *Proceedings of the International Computer Music Conference*, Barcelona, Spain, Sept. 5-9 2005.
- [7] I. Peretz, D. Gaudreau, and A.-M. Bonnel, “Exposure effects on music preference and recognition,” *Memory and Cognition*, vol. 26, no. 5, pp. 884–902, 1998.
- [8] S. Beck, *Making Music with Csound*. Cambridge, MA, USA: MIT Press, January 2000, ch. Acoustic Viability: A Design Aesthetic for Synthetic Musical Instruments.
- [9] E. Maestre, M. Blaauw, J. Bonada, E. Guaus, and A. Perez, “Statistical modeling of bowing control applied to violin sound synthesis,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 4, pp. 855–871, May 2010.
- [10] P. Cook and G. Scavone, “The synthesis toolkit (STK), version 2.1,” in *Proc. Internat’l Computer Music Conf.*, Beijing, China, 1999.
- [11] A. Kontogeorgakopoulos and G. Kouroupetroglou, “Low cost force-feedback haptic interaction with haptic digital audio effects,” in *Lecture Notes in Artificial Intelligence: Gesture and Sign Language in Human-Computer Interaction and Embodied Communication*, E. Efthimiou, G. Kouroupetroglou, and S.-E. Fotinea, Eds., vol. 7206. Springer Verlag, 2012, pp. 48–57.
- [12] —, “Simple cases of low cost force-feedback haptic interaction with haptic digital audio effects,” in *Proc. 9th International Gesture Workshop*, Athens, Greece, 2011.
- [13] A. Kontogeorgakopoulos and C. Cadoz, “Designing and synthesizing delay-based digital audio effects using the CORDIS-ANIMA physical modeling formalism,” in *Proceedings of the Sound and Music Computing Conference*, Berlin, Germany, 2008.
- [14] S. Sastry, *Nonlinear Systems: Analysis, Stability, and Control*. Springer, New York, NY, 2004.
- [15] M. Mathews and J. Smith III, “Methods for synthesizing very high Q parametrically well behaved two pole filters,” in *Proc. Stockholm Musical Acoustic Conference (SMAC)*, Stockholm, Sweden, August 6-9 2003.
- [16] E. Berdahl and J. Smith III, “An introduction to the Synth-A-Modeler compiler: Modular and open-source sound synthesis using physical models,” in *Proceedings of the Linux Audio Conference*, Stanford, CA, USA, April 2012.
- [17] J. Chowning, “The synthesis of complex audio spectra by means of frequency modulation,” *Journal of the Audio Engineering Society*, vol. 21, no. 7, pp. 526–534, September 1973.
- [18] A. Kontogeorgakopoulos, Ph.D. dissertation, Grenoble Institute of Technology, Grenoble, France, 2010.
- [19] E. Berdahl and A. Kontogeorgakopoulos, “The Fire-Fader: Simple, open-source, and reconfigurable haptic force feedback for musicians,” *Computer Music Journal*, vol. 37, no. 1, pp. 23–34, Spring 2013.
- [20] J. Pierce and S. Van Duyne, “A passive nonlinear digital filter design which facilitates physics-based sound synthesis of highly nonlinear musical instruments,” *Journal of the Acoustical Society of America*, vol. 101, no. 2, pp. 1120–1126, February 1997.

A. IMPLEMENTATION OF AN INTERNALLY ENERGY-PRESERVING RESONATOR

For the purposes of shortening the following discussion, the resonator will be assumed to have zero internal damping.⁵ Thus, the energy U stored internally in the spring and mass shown in Figure 1 can be written as the sum of the kinetic energy stored in the moving mass and the potential energy stored in the spring:

$$U = \frac{1}{2}m\dot{x}^2 + \frac{1}{2}kx^2, \quad (3)$$

where x is the position of the mass, \dot{x} is the velocity of the mass, k is the stiffness of the spring, and m is the mass. Assuming that there is very little damping, the resonance frequency

$$f_0 = \frac{1}{2\pi} \sqrt{\frac{k}{m}}. \quad (4)$$

Thus, the resonance frequency can be adjusted by changing k and or m ; however, doing so also will generally affect the potential energy. That is detrimental to the stability—if the resonator is generating energy when it is being frequency modulated, then that energy can cause the vibrations of the resonator to suddenly become very large (i.e. too loud to be conveniently listened to) or this energy can spread to other parts of the model, in this case being imparted to the haptic device and potentially driving it unstable.

In prior work, it was observed that k could be changed without affecting the energy U if k was changed only when

⁵ For the purposes of argumentation, the damping could in any case be implemented externally to the resonator.

$x = 0$ [20]. However, for the fast frequency modulation required to achieve bright timbres, the resonance frequency must be changed much faster than the time required for the resonator to oscillate through half of a carrier cycle; therefore, this approach would not be practical for the present application.

Consequently, an approach was developed for holding the energy U constant subject to an instantaneous change at time τ in both the stiffness and mass. The analysis is performed in the continuous time domain:

$$U = \frac{1}{2}m_{old}\dot{x}^2 + \frac{1}{2}k_{old}x^2 \triangleq \frac{1}{2}m_{new}\dot{x}^2 + \frac{1}{2}k_{new}x^2. \quad (5)$$

x and \dot{x} are required to be the same before and after the change in resonance frequency in order to prevent transients from occurring in the output signal from the resonator (see (5)). Letting the time-varying mass be represented by the function $m(t)$, then $m(\tau_-) = m_{old}$ and $m(\tau_+) = m_{new}$, and similarly for the stiffness function $k(t)$, $k(\tau_-) = k_{old}$ and $k(\tau_+) = k_{new}$. Then, given k_{old} , m_{old} , the previous resonance frequency $f_{old} = \frac{1}{2\pi}\sqrt{\frac{k_{old}}{m_{old}}}$, and the new target resonance frequency f_{new} , the unknowns k_{new} and m_{new} can be calculated.

For instance, by defining $\hat{k}_{new} \triangleq \eta k_{old}$ and $\hat{m}_{new} \triangleq \zeta m_{old}$, then one can arrive at

$$\eta = \frac{\kappa^2(m_{old}\dot{x}^2 + k_{old}x^2)}{m_{old}\dot{x}^2 + \kappa^2k_{old}x^2}, \quad (6)$$

where $\kappa = f_{new}/f_{old}$. ζ can be determined using $\zeta = \eta/\kappa^2$. Under ideal conditions, the derivation would be finished, but to avoid error accumulation in k_{new} and m_{new} over time, it can be a good idea to rescale them to precisely match the resonance frequency despite any numerical round-off errors, etc. Defining $\hat{f}_{new} \triangleq \frac{1}{2\pi}\sqrt{\frac{\hat{k}_{new}}{\hat{m}_{new}}}$, we have the following:

$$k_{new} = \sqrt{\frac{2\pi f_{new}^2}{\hat{f}_{new}} \cdot \hat{k}_{new}} \quad \text{and} \quad (7)$$

$$m_{new} = \sqrt{\frac{\hat{f}_{new}}{2\pi f_{new}^2} \cdot \hat{m}_{new}}. \quad (8)$$

Experimentally it was observed that (7) and (8) appear to prevent the state variables from drifting over time. This approach was employed to control a modified version of the *resonators* object from Synth-A-Modeler [16]. The stiffness and the mass were readjusted at each subsequent audio sample. Thus, the resonator was implemented via a state-space approach using a rotation matrix, where the signals $f(t)$ and $m(t)$ were employed to tune the resonator in real time [15]. It was necessary to provide a seed beginning mass value of $m(0)$. With this approach, $m(t)$ affected the scaling of only the input force into the state variables of the oscillator, and not the recursive computations involving the state variables. As a consequence, the resonator could quite easily tolerate the time-varying parameters without internally becoming unstable.

In testing, it was determined experimentally that this resonator implementation allowed for FM synthesis even with

extreme parameters without the model exploding. This meant that it was quite easy to employ the model without worrying about it creating unpleasant sounds or large forces. In other words, the FM synthesis parameters could be tuned over a large range enabling tangible interaction with an FM synthesis algorithm even for frequency deviations as large as $d = 100.0$ and f_m as large as $100f_c$. From a stability standpoint, it worked considerably better than other more simplistic implementations of the frequency-modulated resonator that had been previously tested by the authors.

B. PRESERVATION OF INTERNAL ENERGY DOES NOT PROVE ASYMPTOTIC STABILITY FOR TIME-VARYING SYSTEMS

It was experimentally determined to be a useful property of the resonator that it preserved its internal energy. However, despite this feature, the model shown in Figure 3 was not always asymptotically stable when the user was touching the resonator (i.e. when the plectrum was behaving like a spring with some damping). This is because, in a time-varying system, even when a component preserves the energy stored in its internal state variables, it still can cause the energy stored in other (even LTI) elements' state variables to grow due to its time-varying impedance.

Consider the case in which a spring with stiffness k_a is permanently attached to the resonator, as depicted in Figure 11. Although the stiffness and mass are updated in order to satisfy (5), the net stiffness of the model, and thus effectively the stiffness of the resonator, is in fact $k(t) + k_a$. Therefore, the total energy of the model

$$U_{total} = \frac{1}{2}m(t)\dot{x}^2 + \frac{1}{2}(k(t) + k_a)x^2 \quad (9)$$

is no longer necessarily preserved. In simulations with a haptic device, it was determined that the resonator tends to self-oscillate at low magnitudes when in contact with the haptic device (see Figures 3 and 10). However, the self-oscillations were small enough that they seemed more of a feature than a bug in a sense that they became part of the sonic and haptic identity associated with the model.




Figure 11. Time-varying resonator with mass $m(t)$ and stiffness $k(t)$ connected to a spring with stiffness k_a .