EFFICIENT MODELING AND SYNTHESIS OF BELL-LIKE SOUNDS

Matti Karjalainen, Vesa Välimäki, and Paulo A. A. Esquef

Helsinki University of Technology
Laboratory of Acoustics and Audio Signal Processing
Espoo, Finland
{firstname.lastname}@hut.fi

ABSTRACT

This paper describes two different techniques that can be used to model and synthesize bell-like sounds. The first one is a source-filter model based on frequency-zooming ARMA (pole-zero) modeling techniques. The frequency-zooming approach is powerful also in modal analysis of bell sound behavior. The second technique is based on a digital waveguide with a single loop filter that is designed to generate inharmonic partials by including one or more second-order allpass sections in the loop filter, possibly augmented with one or a few parallel resonators. A small handbell with inharmonic partials was recorded and used as a target of modeling and synthesis. Sound examples are found in http://www.acoustics.hut.fi/demos/dafx02/.

1. INTRODUCTION

Sound production in bells and bell-like musical instruments has been mastered experimentally for thousands of years. Acoustically we understand it to be composed of decaying sinusoids. After hitting a bell by a hard or a soft object the autonomous vibration consists of eigenmodes that decay exponentially, and this response can be approximated as a linear and time-invariant (LTI) system.

Bells are manufactured in different sizes and shapes, from large and loud bells for churches to small handbells for musical or non-musical purposes. Bells can be tuned, i.e., designed to approach a series of harmonic partials, but it is essential to bell-like sounds that there always are inharmonic components (at least one). This means an irregular spectrum and therefore the efficient synthesis techniques available for harmonic sounds cannot be employed directly.

Another typical property of bells is the warble sound, i.e., clearly noticeable beating of partials, which may be a desirable feature to make the sound more colorful and pleasant to listen to. The warble is caused by minor asymmetry of structure or material so that the modal frequencies in different direction of vibration are slightly different. Unless excited exactly in such a single 'submode' direction, a beating warble sound is produced as a summation of 'submodes'. A good introduction to the acoustics and sound properties of bells is found in [1]. Among other useful sources of information, including sound examples, are [2, 3, 4].

Individual bell-like sounds can be synthesized easily by sampling techniques (i.e., playing back from recordings). However, if the goal is to gain flexibility by parametric control for musical purposes, or to reduce memory requirement of long wavetables, other methods become more attractive. In this paper we discuss two approaches that are computationally efficient and parametrically controllable: (a) modal filterbanks, optionally based on subband techniques, and (b) inharmonic waveguide methods. Both approaches are found interesting also from a more general viewpoint of music-related DSP.

2. SIGNAL ANALYSIS OF BELL SOUND

The exponentially decaying modal response, after the initial attack of hitting a bell, can be utilized in sound synthesis by modeling the radiated sound as a convolution of an excitation signal and an LTI system of modal resonances. This means the applicability of source-filter modeling for sound synthesis. The problem at hand is to estimate a low-order filter which has a good match to the decay part of response and then to find a proper excitation. A strong excitation may cause some nonlinearity in the vibration of a bell, but in most cases the LTI assumption works well.

There are many methods available for estimating a parametric model based on the impulse response of a given LTI system: AR (autoregressive) modeling for all-pole filter and ARMA (autoregressive moving average) modeling for pole-zero filter design [5, 6, 7]. Each eigenmode should correspond to a complex-conjugate pole pair in the resulting filter. Although powerful in many applications, direct application of these methods to bell sounds may turn out to be problematic as will be described below. Very closely located modal frequencies require high-resolution spectral analysis tools, such as the frequency-zooming technique to be discussed.

2.1. Example handbell case

Figure 1 illustrates the temporal and spectral properties of a small handbell the response of which is used as a model target of this study. Subplot 1(a) shows the decay envelope of the bell sound, which deviates from a simple exponential curve, hinting potential beats and warble. This is indeed the case, which is easy to confirm by listening.

Fig. 1(b) depicts the magnitude spectrum of the sound example. At a first glance it may be noticed that there is a set of relatively isolated sharp resonances (frequencies up to 10 kHz are listed in Table 1). A closer look by spectral zooming reveals, however, that most of the resonances have two or even more nearby peaks. Fig. 1(c) shows a zoomed-in magnitude spectrum for the first resonance group around 1313 Hz, showing the existence of two almost equally intense peaks with frequency separation of 2.5 Hz.

Bandpass filtering (bandwidth = 600 Hz) and temporal envelope analysis of the partials up to 10 kHz is plotted in Fig. 2. This reveals clearly the beating in the envelopes of the two lowest par-

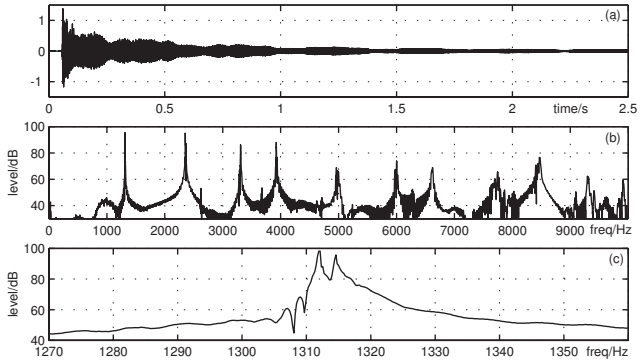

Figure 1: *Analysis and modeling of a small bell sound: (a) recorded time-domain signal, (b) magnitude spectrum up to 10 kHz, (c) magnitude spectrum in the modal region around 1313 Hz.*

Freq/Hz	T60/s
1314.5 / 1312.0	7.04 / 8.02
2362.9 / 2353.3	4.53 / 3.93
3306.5 / 3309.4	1.94 / 0.30
3928.2 / 3923.8	0.57 / 1.77
4993.7 / 4966.6	0.58 / 0.53
5994.4 / 6003.0	0.57 / 0.47
6619.7 / 6598.9	0.14 / 0.04
7671.7 / 7753.2	0.05 / 0.20
8413.1 / 8453.3	0.25 / 0.07
9305.2 / 9292.4	0.17 / 0.12
9912.4 / 9602.3	0.74 / 0.02

Table 1: *Center frequencies of modal pairs (= partials) up to 10 kHz and corresponding effective decay times (T60) in the bell of the case study.*

tials, and some irregularity also in the higher partials. Beating of the partials is a perceptually important feature and must be incorporated properly in synthesis models in order to generate realistic sounding synthesis.

2.2. Frequency-zooming ARMA modeling of modal groups

The modal behavior of individual partials can be analyzed and modeled in various ways. Bandpass filtering around a partial frequency and envelope detection or short-time Fourier analysis to track envelope trajectories can be used and the overall decay time and beating can be approximated by relatively simple rules. For complex modal combinations this may not work, however.

Estimation of a parametric model through AR or ARMA modeling is in principle an attractive way to obtain a digital filter representing the target sound example. It is known, however, that very closely located poles are often too demanding in practical estimation, requiring extremely high filter orders that are impractical and end up to numerical problems. We have developed a technique called *frequency-zooming ARMA* (FZ-ARMA) analysis introduced in [8] that is a modification of known ARMA methods, particularly the Steiglitz-McBride method [7, pp. 174–177]. It is influenced by frequency-selective linear prediction and Prony's method in [9] and [10]. FZ-ARMA estimation and modeling is formulated as follows:

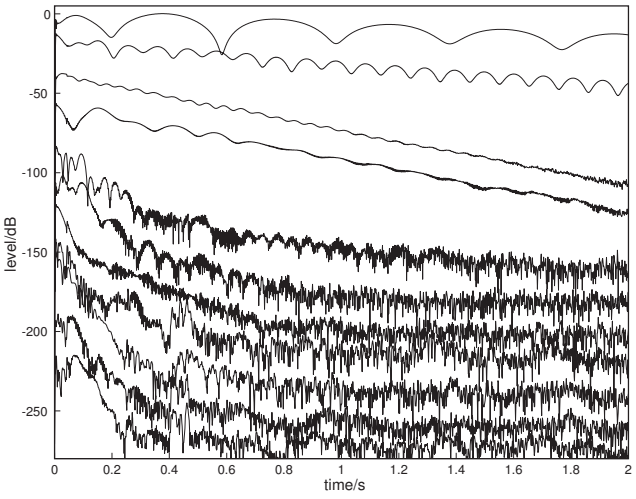

Figure 2: *Decay envelopes for partials of Table 1 from top to down plotted on dB scale with 20 dB offsets for clarity. Higher partials reach the measurement noise floor.*

The FZ-ARMA analysis starts by modulating (heterodyning) the desired frequency band of signal $h(n)$ down to the neighborhood of zero frequency [11, 12] by

$$h_{\mathrm{m}}(n) = e^{j\Omega_{\mathrm{m}} n}\, h(n) \tag{1}$$

where $\Omega_{\mathrm{m}} = 2\pi f_{\mathrm{m}}/f_{\mathrm{s}}$, f_{m} is the modulation frequency, and f_{s} is the sample rate. In the z-domain this can be interpreted as clockwise rotation of poles z_i by angle Ω_{m}, i.e.,

$$\Omega_{i,\mathrm{rot}} = \Omega_i - \Omega_{\mathrm{m}} = \arg(z_i) - \Omega_{\mathrm{m}} \tag{2}$$

but retaining the pole radius. The next step to increase frequency resolution is to limit the frequency range by decimating, i.e., lowpass filtering and down-sampling the rotated response by zooming factor K_{zoom} to obtain a new sampling rate $f_{\mathrm{s,zoom}} = f_{\mathrm{s}}/K_{\mathrm{zoom}}$. This implies mapping to a new z-domain where poles are scaled by the rule

$$z_{i,\mathrm{zoom}} = z_i^{K_{\mathrm{zoom}}} \tag{3}$$

Together mappings (2) and (3) yield new poles [1]

$$\hat{z}_{i,\mathrm{zoom}} = |z_i|^{K_{\mathrm{zoom}}}\, e^{j\,(\arg(z_i)-\Omega_{\mathrm{m}})\,K_{\mathrm{zoom}}} \tag{4}$$

Now it is possible to apply any AR or ARMA modeling to the modulated and decimated response. Notice that this new signal is complex-valued due to the one-sided modulation operation.

The advantage gained by frequency zooming is that in the zoomed subband the order of (ARMA) analysis can be reduced by increasing the zooming factor K_{zoom} and, consequently, the solution of poles and zeros as roots of the denominator and numerator polynomials of the model transfer function is simplified. Additionally this means that a different resolution can be used in each subband, for example based on knowledge about modal complexity and perceptual relevance of a subband.

[1] Note that Eqs. (2) and (3) merely characterize how the z-domain properties of a given response are changed through modulation and decimation, but the estimated pole-zero pattern of an FZ-ARMA model will be obtained only in the next step.

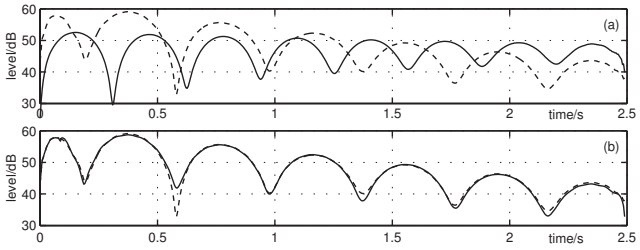

Figure 3: *Modeling of the decay envelope of the first partial (see Fig. 2 top curve) with FZ-ARMA orders: (a) ARMA(0,4) and (b) ARMA(40,4). Dashed line is the target and solid line the model behavior.*

After solving the poles within a zoomed subband they can be remapped to the full sample rate by inverse scaling of the radii of poles as well as rotating them counter-clockwise, i.e.,

$$\hat{z}_i = \hat{z}_{i,\mathrm{zoom}}^{\,1/K_{\mathrm{zoom}}}\, e^{j\Omega_{\mathrm{m}}} \tag{5}$$

Because of the one-sided down-modulation used in (1), each pole z_i must be used as a complex conjugate pair in order to obtain real-valued filters.

Finally there are two alternatives of complete model construction. The subband filters, mapped back to full sample rate and appropriate frequency band, can be combined into a composite filter, such as a parallel filterbank. If the computational complexity of the composite filter is too high and if the subbands are uniform or regular in bandwidth, multirate techniques [13] can be utilized, although this introduces extra processing latency.

In the investigations below the frequency-zooming method used for solving FZ-ARMA coefficients is the Steiglitz-McBride iteration (`stmcb` in MATLAB). Notice that the IIR filter orders N for numerator and P for denominator refer to real-valued filters, with complex-conjugate pairs constructed from one-sided zeros and poles obtained from the model of the decimated signal. Thus the orders of real-valued filters are twice the numbers of zeros and poles from the above procedures, respectively.

An FZ-ARMA denominator of order 4 is a natural choice for a partial with regular beating, while order 2 is sufficient for simple exponential decay. In practice order 6 may be needed to account for a bit irregular beating decay. The order of the numerator depends on how well the transient attack needs to be modeled, and will be discussed later below.

Figure 3 illustrates how the FZ-ARMA analysis works for the first partial of the bell sound characterized in Fig. 1. In Fig. 3(a) on approximation of the target response using a zero/pole model of order (0,4) yields notable deviation in the decay envelope. Order (40,4) in (b) shows a very good fit. The same can be achieved by order (4,6), i.e., with more poles and much less zeros. A zooming factor K_{zoom} of 200 was applied.

2.3. Real-time synthesis of FZ-ARMA bell models

The target bell sound was modeled for excitation-filter synthesis in the following way. First the bulk delay in the recorded target sound was removed. Frequencies above 10 kHz did not contribute to the perception of recorded sound noticeably, so the modeling was applied only to a lowpass-filtered signal for modes shown in Table 1. For all partials, it was sufficient to set the order of the

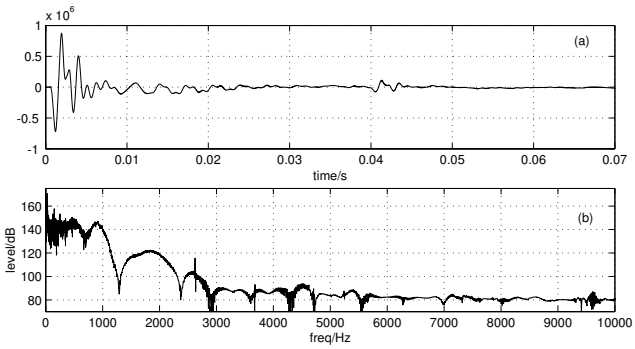

Figure 4: *Properties of the residual used in source-filter synthesis: (a) 70 ms from the beginning of the residual and (b) spectrum of the whole residual.*

AR part of the FZ-ARMA filter equal to 4. As for the MA part, orders between 10 and 30 were enough to properly model the intial transient. This FZ-ARMA analysis was actually used to estimate the modal frequencies and decay times in Table 1.

The poles obtained from the FZ-ARMA subband models were combined to form a single all-pole synthesis filter. The target signal was then inverse filtered with the synthesis filter in order to obtain an excitation for perfect synthesis. Figure 4 depicts the first 70 ms of the excitation and the spectrum of the whole excitation. Note that the FZ-ARMA had a slight tendency to overestimate the pole radii, which leads to antiresonances in the residual spectrum at some modal frequencies.

The excitation was truncated, with end tapering to avoid discontinuity, and was used as an input signal for source-filter synthesis. An interesting detail is a secondary excitation around 40 ms, probably due to double impact in hitting the bell. If the excitation wavetable is made shorter than 40 ms, this detail is lost in the synthesized sound. (In fact, there are even more minor extra excitations in the target sound later in time, as can be noticed by careful listening to the original sound.)

Now we have obtained a highly efficient synthesis procedure which yields an almost perfect resynthesis, only the later excitations as well as the recording noise of the target sound are missing. A 60 ms wavetable corresponds to about 2600 samples for a sample rate of 44100 Hz, and the order of the obtained all-pole synthesis filter is $4 \times 11 = 44$. The filter order can be reduced further without audible difference by using four poles only for the first to partials exhibiting systematic beating and just a pole pair for others, thus ending up with a filter order of $4 \times 2 + 2 \times 9 = 26$. In addition to the low computational load the system is now parametrized, i.e., the filter parameters can be changed for different sounds or even modulated for interesting audio effects. The excitation can be selected among different variants from a collection of wavetables. Multiple impacts can be simulated just by reading multiple wavetables with offset between their triggering time.

3. INHARMONIC DIGITAL WAVEGUIDE SYNTHESIS

A very efficient way to synthesize harmonic sounds of exponentially decaying partials is by using digital waveguides [14]. The regularity of harmonic spectral structure can be maximally exploited by a feedback loop of a (possibly fractional) delay and a lowpass loop filter. In commuted synthesis the harmonic resonator is modeled by a digital waveguide and everything else in sound

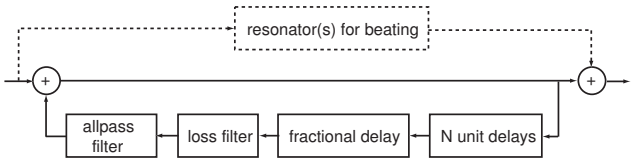

Figure 5: *Inharmonic waveguide model with strong dispersion. Parallel resonator(s) for additional inharmonicity and beating of partials shown in dashed line.*

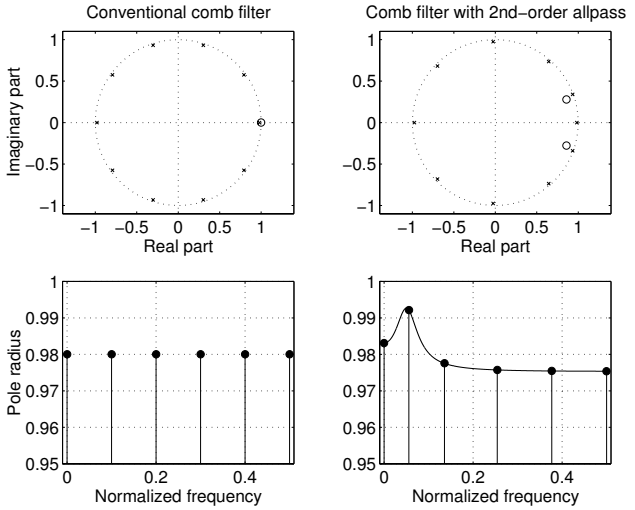

Figure 6: *(top) Poles and zeros of a conventional comb filter (left) and the one with a nested second-order allpass filter (right). The corresponding pole radii are shown below as function of frequency. The solid line (bottom right) is the loop gain determined with the group delay of the open loop (see Eq. 8).*

synthesis is consolidated in a wavetable for excitation of the filter. An advanced version of such single loop filter model for the acoustic guitar is described in [15].

Since the bell sounds are always inharmonic, the regularity of the basic digital waveguide cannot be exploited directly. A digital waveguide with smooth dispersion has been used for dispersive piano string models [16, 17]. A banded waveguide structure has been suggested for producing inharmonic spectra for vibrating bars [18]. Inharmonic comb filters with an allpass filter in the feedback loop have been used for generating spectra of spherical resonators [19]. We will use a similar structure and utilize the resonance characteristics of a second-order allpass filter. Figure 5 depicts a diagram of a digital waveguide, including a bulk delay, a fractional delay for pitch adjustment, a low-order low-pass filter for frequency-dependent decay-time control of partials, and an inharmonizing second-order allpass filter. There can be additional resonators in parallel with this (as shown in dashed line) for additional inharmonicity or controlling the beating of partials [20].

3.1. Comb filter with a nested second-order allpass filter

Inserting a second-order allpass filter into the loop of a comb filter enables design of very inharmonic comb filters. The locations of resonances depend on the change of phase in the feedback loop. The phase response of the second-order allpass is [21]

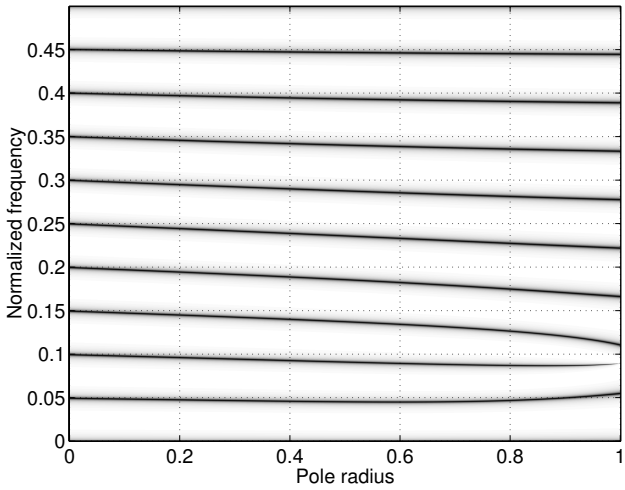

Figure 7: *Magnitude response of the system in grey scales (white is 0 and black is the maximum value) as a function of pole radius of the allpass filter ($L = 18$, $g = 0.995^{10}$).*

$$\begin{aligned}\phi_A(\omega) = &-2\omega - 2\arctan\left[\frac{r\sin(\omega-\theta)}{1-r\cos(\omega-\theta)}\right] \\ &- 2\arctan\left[\frac{r\sin(\omega+\theta)}{1-r\cos(\omega+\theta)}\right]\end{aligned} \tag{6}$$

where ω is the normalized frequency and r and θ are the radius and frequency of the pole of the second-order allpass filter, respectively. The total phase function of the feedback loop (open loop) depends on both the allpass filter and the delay line, which contributes a linear phase term $-L\omega$, where L is the delay-line length:

$$\begin{aligned}\phi_{\text{loop}}(\omega) = &-(L+2)\omega - 2\arctan\left[\frac{r\sin(\omega-\theta)}{1-r\cos(\omega-\theta)}\right] \\ &- 2\arctan\left[\frac{r\sin(\omega+\theta)}{1-r\cos(\omega+\theta)}\right]\end{aligned} \tag{7}$$

Resonances occur at frequencies where the value of ϕ_{loop} is a multiple of 2π.

The decay time of the resonances is also warped by the allpass filter. The pole radii of a conventional comb filter are $g^{1/L}$, where g is the loop gain. The pole radii of the modified system depend on the group delay of the open loop:

$$|p_k| = g^{1/\tau(\omega_{p,k})} \tag{8}$$

where $\tau(\omega_{p,k})$ are the open-loop group delay values at pole frequencies $\omega_{p,k}$. The group delay depends on both the delay-line length L and the allpass filter, as it is the negative derivative of (7).

Figure 6 shows the pole-zero diagrams of a conventional comb filter ($L = 10$) and the system of interest in this paper (delay line with $L = 8$ and a second-order allpass filter). It is seen that while the poles are uniformly spaced in the conventional comb filter, this is not the case for the comb filter with a second-order allpass filter. In the bottom right in Fig. 6, the dependence of the pole radii on the group delay of the allpass filter according to (8) is demonstrated.

The selection of the allpass filter parameters is not as easy as it may seem, because they affect the properties of the system in a

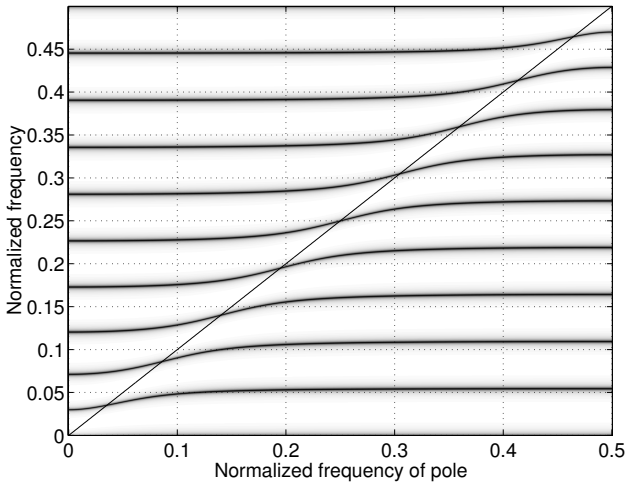

Figure 8: *Magnitude response of the system in grey scales (white is 0 and black is the maximum value) as a function of pole frequency of the allpass filter. The solid line shows the location of the pole ($L = 18$, $g = 0.995^{10}$).*

complicated way. Figures 7 and 8 demonstrate how changing the value of pole radius R or pole frequency θ affects the locations of resonances, when one of the parameters is fixed (θ = 0.08 in Fig. 7 and R = 0.99 in Fig. 8). The delay-line length L is 18 samples. The resonances locations of a normal comb filter are the same as those in the left end of Fig. 7 (obtained with R = 0). It is seen that the presence of an allpass filter moves all resonances, except for the trivial ones at 0 and the Nyquist frequency.

3.2. How to design an inharmonic waveguide

Since there are only two degrees of freedom in the allpass filter, it is reasonable to use them to control the two lowest resonance frequencies of the system, which are usually the perceptually most important ones. We need another loop filter to control the decay times. This loss filter may be a linear-phase filter, such as a 3-tap FIR filter, so that it does not affect the inharmonicity of the system. Then it is possible to first design the allpass filter and to account for its effect on the decay time of modes when designing the loss filter.

As the low-frequency partials of the bell sound are most important from a perceptual point of view, we tried to match only the first 3 partials. The first partial also plays an important role in the perception of the pitch of the bell sound. Therefore, as a starting point, we decided to set the length of the delay-line and the value of the fractional-delay for the tuning filter as to have a harmonic tone with the fundamental matching the first partial of the bell tone. From the values shown in Table 1, we verify that the frequencies of second and third partials of the tone are lower than those of the harmonic one. The task of the inharmonizing allpass filter then will be to attract the second and third partials towards the first partial.

The capability of the allpass filter to attract neighbor resonances is clearly seen in Figs. 7 and 8. From Fig. 7 we also verify that the attraction power increases with the pole radius. In our case, it seems natural to set the frequency of the allpass pole between the frequencies of the first and third partials. In fact, for optimization

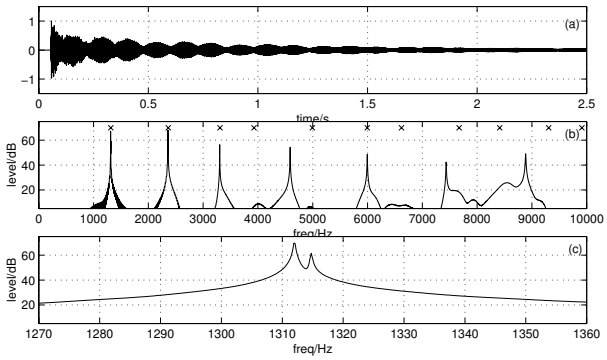

Figure 9: *Analysis of a synthetic bell sound: (a) time-domain signal, (b) magnitude spectrum up to 10 kHz with crosses indicating the target frequencies of the partials, and (c) magnitude spectrum around the first partial showing the pair of peaks.*

purposes, we can relax about matching the precise values of the partial frequencies and care more about their ratios. This allows us to play only with 2 parameters: the radius and frequency of the allpass pole. Once the frequencies of the first 3 partials are related by the desired ratios the matching to the actual partial frequencies can be achieved by adjusting the length of the delayline and the parameter of the tuning filter.

For the tuning filter we used a first order allpass filter the parameter of which is given by [22]

$$a_t = \frac{\sin[(1-\delta)\omega_1/2]}{\sin[(1+\delta)\omega_1/2]}, \tag{9}$$

where ω_1 is the frequency of the first partial in radians and $0 < \delta < 1$ is the fractional part of the delay line.

We tuned an inharmonic comb filter to match the 3 lowest modes of the handbell sound. With a delay line of $L = 28$ plus a fractional delay of $\delta = 0.995$ and allpass filter with parameters $\theta = 2495/f_s$ and $R = 0.913$ we obtain a reasonably good match with the 3 resonances as depicted in Fig. 9. Altogether the structure produces 7 inharmonic modes up to 10 kHz. This test case shows that it is possible to synthesize bell-like sounds that are in tune and sound similar to the target sound. Note that the sixth partial was also closely matched, although unintentionally.

The computer-assisted design process for this test case required some time and nerves, and obviously an automatic design algorithm is needed. It may either use an optimization strategy, such as a genetic algorithm, or a heuristic search based on what we know and can learn from the properties of the structure.

The choice of a 3-tap FIR filter for the loss filter allows us to meet the magnitude response requirements at two frequencies precisely. The main restriction to this approach is that the target magnitude at the higher frequency must be less than that of the lower one, in order to obtain a loss filter with a lowpass characteristic. Again, for psychoacoustic reasons we decided to set the magnitude response of the loss filter to match the decay times of the first two partials. Moreover, one has to remember that the presence of the allpass filter in the feedback loop affects the decay time of modes, and this effect has to be taken into consideration when defining the specification for the loss filter.

If the frequencies of the first two partials in radians are denoted by ω_1 and ω_2, the magnitude response of the loss filter at ω_1 and

ω_2 is given by

$$g(\omega_k) = e^{\tau(\omega_k)\log(0.001)/(T60_k f_s)}, \quad k = 1 \text{ and } 2, \qquad (10)$$

where $\tau(\omega_k)$ are the measured group delays of the open loop at ω_k and $T60_k$ are decay times of the first and second partials, as indicated in Table 1.

Although a single inharmonic comb filter suffices to replicate an inharmonic tone, it does not produce the prominent beating that is perceived in the lowest partials of a bell tone. To imitate this effect we can add a second inharmonic comb filter with parameters slightly changed to approximate the other resonant modes of the first three partials. Alternatively, we can match these modes by adding second-order resonators in parallel with the inharmonic comb filter.

The transfer function of a second-order resonator can be written as

$$H_r(z) = \frac{b_0}{1 - (2R\cos(2\pi f_r/f_s))z^{-1} + R^2 z^{-2}}, \qquad (11)$$

where f_r is the frequency of the resonance in Hz, f_s is the sample rate, b_0 is a gain factor, and R is the radius of the poles. The value of R can be derived from the T60 value by

$$R = e^{\log(0.001)/(T60 f_s)}. \qquad (12)$$

In our example, we used two extra resonators to produce the beating effect in the first and second partials. Thus, one resonator was tuned at $f_r = 1312.0$ Hz and R was set according to Eq. (12) with $\mathrm{T60} = 8.02$ s. For the other resonator we used $f_r = 2353.3$ Hz and $\mathrm{T60} = 3.93$ s. The magnitude spectrum zoomed in around the first partial of the synthetic tone is shown in Fig. 9(c), where two peaks with a frequency difference of 2.5 Hz are seen.

4. ACKNOWLEDGMENT

The work of Paulo Esquef has been financed by the Brazilian National Council for Scientific and Technological Development (CNPq) and the Academy of Finland, project "Technology for Audio and Speech Processing".

5. REFERENCES

[1] N. H. Fletcher and T. D. Rossing, *The Physics of Musical Instruments*. New York, USA: Springer-Verlag, 1991.

[2] T. D. Rossing, *Acoustics of Bells*. Striudsburg, PA: Van Nostrand-Reinhold, 1984.

[3] B. Hibberts, "The Sound of Bells." http://www.hibberts.co.uk/.

[4] J. Ketteringham, "Library of Bell Recordings." http://www.eccentrix.com/personal/johnketteringham/.

[5] S. M. Kay, *Fundamentals of Statistical Signal Processing: Volume I: Estimation Theory*. Englewood Cliffs, N.J.: Prentice-Hall, 1993.

[6] A. V. Oppenheim, A. Willsky, and I. Young, *Signals and Systems*. Englewood Cliffs, New Jersey: Prentice-Hall, 1983.

[7] M. H. Hayes, *Statistical Digital Signal Processing and Modeling*. John Wiley & Sons, 1996.

[8] M. Karjalainen, P. A. A. Esquef, P. Antsalo, A. Mäkivirta, and V. Välimäki, "AR/ARMA Modeling of Modes in Resonant and Reverberant Systems," in *Preprint 5590 AES 112th Convention*, (Munich), 2002 May.

[9] J. Makhoul, "Spectral Linear Prediction: Properties and Applications," *IEEE Trans. ASSP*, vol. 23, no. 5, pp. 283–296, 1975.

[10] J. Laroche, "A New Analysis/Synthesis System of Musical Signals Using Prony's Method – Application to Heavily Damped Percussive Sounds," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, pp. 2053–2056, vol. 3, 1989.

[11] J. A. Moorer, "Signal Processing Aspects of Computer Music: A Survey," *Proc. IEEE*, vol. 65, pp. 1108–1137, Aug. 1977.

[12] V. Välimäki and T. Tolonen, "Development and Calibration of a Guitar Synthesizer," *J. Audio Eng. Soc.*, vol. 46, no. 9, pp. 766–778, Sept. 1998.

[13] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs, NJ: Prentice Hall, 1993.

[14] J. O. Smith, *Principles of Digital Waveguide Models of Musical Instruments*, ch. 10 in *Applications of Digital Signal Processing to Audio and Acoustics*, (ed. Kahrs and Brandenburg). Kluwer Academic Publishers, 1998.

[15] M. Karjalainen, V. Välimäki, and T. Tolonen, "Plucked-String Models, from the Karplus-Strong Algorithm to Digital Waveguides and Beyond," *Computer Music Journal*, vol. 22, no. 3, pp. 17–32, 1998.

[16] S. A. V. Duyne and J. O. Smith, "A Simplified Approach to Modeling Dispersion Caused by Stiffness in Strings and Plates," in *Proc. Int. Computer Music Conf.*, (Aarhus, Denmark), pp. 407–410, Sept. 1994.

[17] D. Rocchesso and F. Scalcon, "Accurate Dispersion Simulation for Piano Strings," in *Proc. Nordic Acoustical Meeting*, (Helsinki, Finland), pp. 407–414, June 1996.

[18] G. Essl and P. R. Cook, "Measurements and Efficient Simulations of Bowed Bars," *J. Acoust. Soc. Am.*, vol. 108, pp. 379–388, July 2000.

[19] D. Rocchesso and P. Dutilleux, "Generalization of a 3-D Acoustic Resonator Model for the Simulation of Spherical Enclosures," *EURASIP J. Applied Signal Processing*, vol. 1, no. 1, pp. 15–26, 2001.

[20] B. Bank, V. Välimäki, L. Sujbert, and M. Karjalainen, "Efficient Physics-Based Sound Synthesis of the Piano Using DSP Methods," in *Proc. European Signal Processing Conf. (EUSIPCO 2000)*, vol. 4, (Tampere, Finland), pp. 2225–2228, Sept. 2000.

[21] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing*, pp. 274–279. Prentice-Hall, second ed., 1999.

[22] K. Steiglitz, *A Digital Signal Processing Primer – with Applications to Digital Audio and Computer Music*, ch. 6. Addison-Wesley Publishing Company, Inc., 1996.