Audio Engineering Society
Convention Paper

Presented at the 117th Convention
2004 October 28–31 San Francisco, California

Interpolating Linear- and Log-Sampled Convolution

D. B. (Don) Keele, Jr. *AES Fellow*

Harman International
Northridge, CA 91329, USA
E-mail: DKeele@Harman.com

ABSTRACT

This paper describes a class of FIR filter/convolvers based on interpolation that allow sparse specification of the filter's impulse-response waveform or equivalently its frequency spectrum in both linear- and log-spaced domains. Interpolation allows the filter's impulse response or frequency response to be specified in significantly fewer samples. This is turn means that far-less filter taps are required. Linear- and log-sampled interpolating filter/convolvers can further be categorized into two types: Type 1, interpolation in time, and Type 2, interpolation in frequency. Type 1 provides direct specification of the filter's impulse response in linear or log time, while Type 2 allows direct specification of the complex (real-imaginary) frequency response of the filter in linear or log frequency. Each form of filter vastly reduces the number of filter taps but greatly increases the processing complexity at each tap. Efficient implementations of the log-spaced filter-convolvers are presented which use multiple asynchronous sample-rate converters. This paper is a continuation of the author's "Log Sampling" paper presented to the AES in Nov. 1994. This paper represents work in progress with a conceptual description of the convolution technique with minimal mathematical development.

0. INTRODUCTION

This paper defines a class of FIR filters and convolvers based on interpolation that allow sparse specification of the filter's impulse time waveform or equivalent frequency spectrum in both linear- and log-spaced domains. Traditionally, taps on an FIR filter are arranged sequentially in a sequence of equal-spaced time delays equal to the sample-rate interval. The outputs of the taps are summed and then feed to a brick-wall low-pass filter forming a sinc-function-based interpolation reconstruction filter. Effectively, each tap has its own sinc reconstruction filter which is spaced in time by the sample rate of the digital processing system. The interpolating filter/convolvers described here have fewer taps that are spaced farther in time than the basic high-speed sample rate of the system. These taps can be linearly or logarithmically spaced.

Linear- and log-sampled interpolating filter/convolvers can further be categorized into two types: Type 1, interpolation in time, and Type 2, interpolation in frequency. The Type 1 interpolation-in-time FIR filter allows specification of the time waveform by a series of sparsely sampled values linearly- or logarithmically-spaced in time. The Type 2 interpolation-in-frequency FIR filter allows specification of the filter's complex (real-imaginary or mag/phase) frequency response by series of sparsely sampled values linearly- or logarithmically-spaced in frequency. All the filters exhibit a one-to-one correspondence between the filter tap weights and the data values in either time or frequency for both the linear and log domains.

The time or frequency sample values themselves can be straightforwardly computed using an alternate definition of the sampling and reconstruction process based on the orthogonal characteristics of the sinc function [1, Section 2.1.2]. This process essentially combines the filtering and sampling of the conventional method in one operation. This process is reviwed in Appendix 1 and can be applied to calculating either linear or logarithmic samples in either time or frequency.

Logarithmic sampling refers to a non-linear impulse response sampling technique that has a high initial sample rate which then falls inversely with time as time progresses [1]. The impulse response of real-world physical systems decays faster at high frequencies than low frequencies. This is due to the approximate constant-Q behavior of the resonators that often make up these systems. Logarithmic sampling is optimum for sampling the impulse response of these types of systems. Other authors have described non-linear sampling techniques that have a high initial sampling rate which then falls with frequency[2], [3]. Appendix 2 outlines the mechanics of log sampling.

The four different kinds of interpolating filters are described in the following sections. The analysis and most of the examples assume continuous-time operation to simplify the concepts. Section 1 describes type 1 linear-sampled interpolation-in-time convolvers. Section 2 describes type 2 linear-sampled interpolation-in-frequency convolvers. Section 3 describes type 1 log-sampled interpolation-in-time convolvers. Section 4 describes type 2 log-sampled interpolation-in-frequency convolvers. Section 5 explores two efficient implementations of the log convolvers using multi-rate techniques that use multiple asynchronous sample rate converters, while section 6 concludes. Appendix 1 reviews an alternate view of sampling, while appendix 2 reviews the concepts of log-sampling in time.

1. LINEAR-SAMPLED TYPE 1: INTERPOLATION IN TIME

1.1. Description

The linearly-sampled Type 1 interpolation-in-time filter is closest to the traditional FIR filter. Here the taps are spaced farther apart in time, effectively forming a lower sample-rate channel which allows the direct specification of a low-frequency waveform with many fewer samples. In this situation, the impulse response of each tap is linearly-sampled interpolation function such as a windowed sinc which is described in Appendix 1 [1, 4, 5] or a truncated cubic spline [1, 6, 7]. In the frequency domain, the interpolation function provides a linear-phase sharp-cutoff low-pass filter. Note however, that the cutoff frequency of each low-pass filter is arbitrarily less than the Nyquist half-sample rate of the DSP hardware that is running the filter.

Fig. 1 shows the block diagram of the linear-sampled type 1 interpolation-in-time convolver. Each coefficient or weight a_0 to a_n sets the level of the corresponding sample at equally-spaced times of t_0 to t_n. Note that the sample times can be arbitrarily lower than the sample times of the hardware that the filter is implemented on. The linear-sampled type 1 filter is closest to the traditional FIR filter where the tap times are much closer together because of the inherent higher sampling rate of the DSP hardware. The linearly-sample type 1 filter allows low-frequency waveforms to be specified with far fewer samples because of the effective lower sample rate of the convolver.

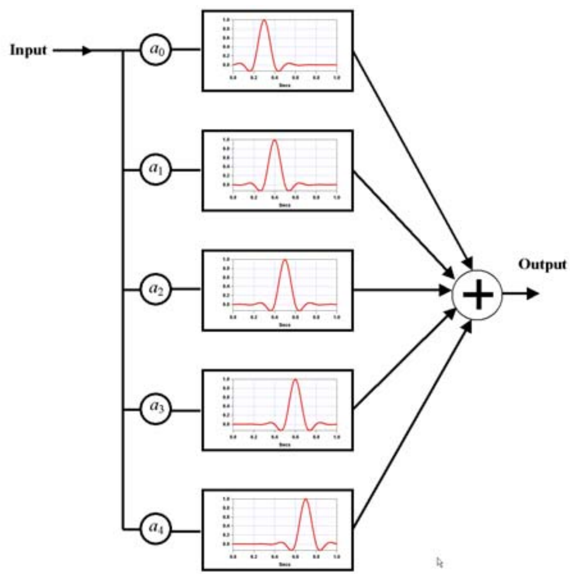

Fig. 1. Block diagram of a linear-sampled type 1 interpolation-in-time FIR filter/convolver. This form of convolver allows the filter's impulse response to be directly specified in linear time. The impulse response of each tap is a linearly-sampled sinc-based interpolation function. Each interpolation function is linearly spaced in time, but at an effective sampling rate which is much less than the sampling rate of the DSP hardware that implements the filter. This allows many fewer samples to specify the waveform of the impulse response of the resultant filter. Each channel is illustrated with a one second time-span convolver containing an interpolator waveform that interpolates the sample values at an effective sample rate of 10 Hz. Five channels are illustrated with linear-spaced sample times ranging from 0.3 to 0.7 s with a step size of 0.1 s.

The convolver output is given by the following equation, which assumes continuous time operation for simplification:

$$y(t) = \sum_{n=0}^{N-1} a_n x(t) \otimes H_{\text{interp}}(t - nT) \qquad (1)$$

where $x(t)$ is the input signal, $y(t)$ is the system output signal, N is the number of channels, the a_n are the amplitude coefficients of each channel, H_{interp} is the interpolation function, T is the sample time of the structure (much lower than the sampling-rate of the DSP hardware), and $\otimes$ is the convolution operator.

Eq. 1.1 states that the output is formed by the sum of the inputs convolved with successively delayed versions of the interpolator function. No surprises here! This is exactly how a standard FIR filter works except that the interpolation functions are effectively spaced at the DSP sampling rate and only one interpolation or reconstruction filter is required at the filters output after the summing block. The convolution operations described in Fig. 1 can alternately be implemented as a conventional FIR filter with only one reconstruction filter operating at a reduced sample rate using decimation.

2. LINEAR-SAMPLED TYPE 2: INTERPOLATION IN FREQUENCY

2.1. Description

The linearly-sampled Type 2 interpolation-in-frequency filter in contrast, allows direct specification of the real and imaginary parts of the filter's frequency spectrum in linear frequency. Here, the frequency spectrum can be sampled much more sparsely than is typical with the DSP hardware. Effectively, the frequency response of the filter is slowly varying in frequency due to the sparseness of the samples. Each linearly-interpolated sample in the filter's frequency response, corresponds to a FIR filter tap whose impulse response is a tone burst at the frequency of the spectrum sample. The waveform of the burst is the inverse Fourier transform of the frequency-domain interpolation function. In this linear case, the burst time of all the filter taps is equal, only the number of cycles in the burst changes. Fig. 2 illustrates this behavior.

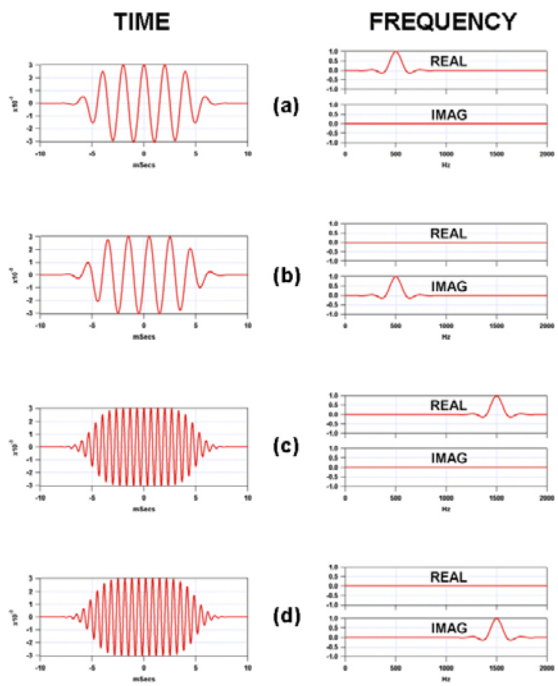

Fig. 2. Illustration of the time response of a cubic-spline interpolator in linear frequency. Note *linear* frequency scale of 0 to 2000 Hz (right), and *linear* time scale of –10 to +10 ms (left). The frequency sample rate is 100 Hz, i.e. the interpolator steps in frequency every 100 Hz. (a) 500 Hz interpolator in real part of spectrum. (b) 500 Hz interpolator in imaginary part of spectrum. (c) 1,500 Hz interpolator in real part of spectrum. (b) 1,500 Hz interpolator in imaginary part of spectrum. Note that the time burst width is constant, only the frequency and phase of the burst changes.

The impulse response of each tap is a constant-frequency tone burst of arbitrary phase (complex) which is the time waveform corresponding to the interpolator at that specific burst frequency. The frequency of each tone burst at a particular tap is incremented at a step size which depends inversely on the duration of the burst. The tone burst is complex (real/imaginary or cosine/sine) because two interpolators exist in the frequency domain, one for the real part of the frequency response and one for the imaginary part. Strictly speaking, causality demands that not only do these two interpolators exist in the real and imaginary parts of the frequency spectrum but their Hilbert transformed counterparts exist in the opposite part of the complex response.

Fig. 3 shows the block diagram of the linear-sampled type 2 interpolation-in-frequency convolver. Each coefficient a_0 to a_n sets the level of the corresponding tone burst associated with a particular interpolation function in frequency. The frequency interpolators are spaced equally in frequency at arbitrary values from f_0 to f_n. Note that the frequency interval can be arbitrarily lower than the sampling frequency of the hardware that the filter is implemented on.

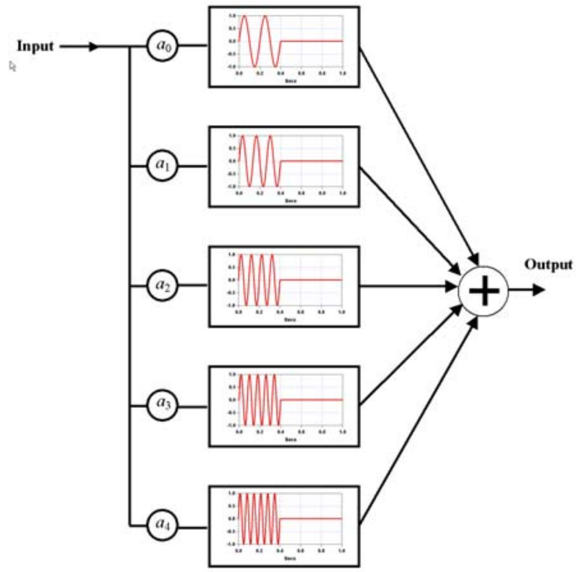

Fig. 3. Block diagram of a linear-sampled type 2 interpolation-in-frequency FIR filter/convolver. Each tap's impulse response is a constant-width tone burst. The frequency and phase of each burst depends on the location of the interpolator in the complex frequency response. This form of convolver allows the filter's frequency response to be directly specified at a relatively small number of linear-spaced points in frequency. Each channel is illustrated with a one second time-span convolver containing a 0.4 s constant-width tone burst of varying frequency. Five channels are illustrated covering a burst frequency range of 5 to 15 Hz with a step size of 2.5 Hz (= 1/BurstWidth = 1/0.4, the sample interval of the frequency interpolator).

For illustration purposes, each burst is shown un-windowed. Practically, the actual bursts will exhibit an amplitude ramp up at the start and a ramp down at the end whose envelope shape depends on the chosen frequency-domain interpolator (see left column in Fig. 2). Alternately, the bursts van be formed by windowing a sine wave (of arbitrary phase) with a Tukey window [5, p. 56] which provides a half-Hann ramp-up and ramp-down to the start and end of the burst respectively.

The linearly-sampled type 2 filter allows complex (real/imaginary or magnitude/phase) frequency responses to be directly specified in the frequency domain at linear-spaced points with far fewer samples than required by conventional FIR filters. The frequency, magnitude, and phase of the tone burst maps directly to the corresponding point in the frequency domain.

The convolver output is given by the following, which assumes continuous time operation for simplification:

$$y(t)=\sum_{n=0}^{N-1} a_n x(t)\otimes\left[\Im^{-1}\left(H_{\text{interp}}\left(f_0-n\Delta f\right)\right)\right] \quad (2)$$

where $x(t)$ is the input signal, $y(t)$ is the system output signal, N is the number of channels, the a_n are the amplitude coefficients of each channel, H_{interp} is the frequency interpolation function, f_0 is the start frequency (can be an arbitrary value and need not be zero), ∇f is the sample frequency interval, $\Im^{-1}$ is the inverse Fourier transform operator, and $\otimes$ is the convolution operator.

This type of convolver is conceptually similar to the previous convolver but each channel contains the time waveform of an *inverse* Fourier transformed interpolation function in frequency. Here the time waveforms are essentially tone bursts of constant width of varying frequency and phase.

As an implementation example, a ten-tap FIR filter can be created that allows linear specification of a 0-to-20 kHz pass-band filter's frequency response in only ten equally-spaced 2 kHz spaced samples.

3. LOG-SAMPLED TYPE 1: INTERPOLATION IN TIME

3.1. Description

The logarithmically-sampled Type 1 filter allows direct specification of the impulse response of the filter in log time. Effectively, the filter's sample rate is high at short times and low at long times. The taps of this filter are spaced logarithmically in time at a sample rate which falls inversely with frequency.

The impulse response of each of the filter's taps is a log-warped sinc function that is shifted up and down in log time. Log sampling in time has been shown to be a very efficient means of storing the samples of a real-world device's impulse response [1]. Log sampling has a high sample rate initially, that then falls inversely with time slowing as time progresses [1] , i.e. $f_S = k/t$ where f_S is the sample rate, t is the time from the start of the impulse, and k is a constant. The effective smoothing that log sampling and reconstruction provides has also been called "complex smoothing" [2]. An efficient convolver implementing the log-sampled Type 1 processing, which makes heavy use of multiple asynchronous sample-rate converters, is described in [1, Sec. 8.1.3 and Fig. 22] and is described again here in Section 6.

Fig. 4 shows the block diagram of the log-sampled type 1 interpolation-in-time convolver.

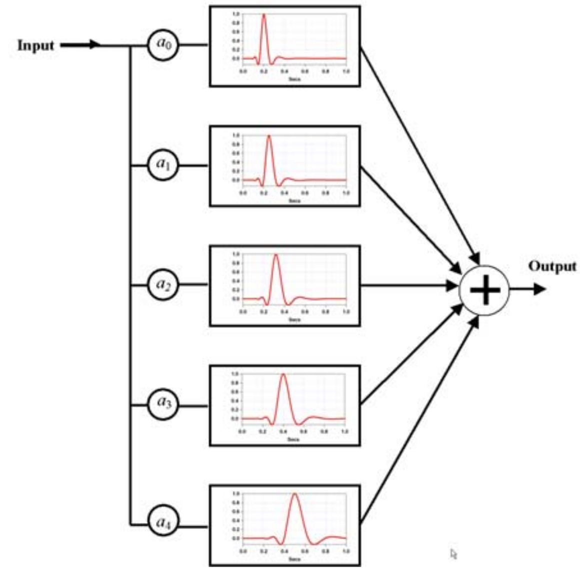

Fig. 4. Block diagram of a log-sampled type 1 interpolation-in-time FIR filter/convolver with five sampling channels. This form of convolver allows the filter's impulse response to be directly specified in log time. The impulse response of each tap is a log-warped sinc-based interpolation function. Each interpolation function is log spaced in time. The log sample points are spaced closer together at short times and farther apart at long times. This allows many fewer samples to specify the waveform of a wide-band impulse that has its high-frequency energy content at short times and its low-frequency energy content at long times. Each channel is illustrated with a one second time-span convolver containing a linearly-sampled log-warped interpolator waveform that interpolates the sample values at an effective sample rate of 10 points per decade of time or a sample ratio of $10^{0.1}$ = 1.2589. Five channels are illustrated with approximate log-spaced sample times of 0.2, 0.25, 0.316, 0.4, and 0.5 s.

In continuous time, the convolver output is given by the following:

$$y(t)=\sum_{n=0}^{N-1} a_n x(t)\otimes H_{\text{interp}}\left(N_e \ln\left(\frac{t}{r^n t_{\min}}\right)\right) \quad (3)$$

where $x(t)$ is the input signal, $y(t)$ is the system output signal, N is the number of channels, the a_n are the amplitude coefficients of each channel, H_{interp} is the interpolation function, N_e is the sample density, r is the ratio between successive sample times ($r = e^{1/N_e}$), $t_{\min}$ is the sampling start time ($t_{\min} > 0$), and $\otimes$ is the convolution operator. Here the sampling stops (last time sample) at $t_{\max} = r^{N-1} t_{\min}$.

This convolver also has the same configuration as Fig. 1, but the log-warped interpolation function is shifted up in time at constant ratio increments from t_{min} to t_{max} in a geometric sequence.

As an example, the impulse response waveform of a wide-band signal can be replicated using far fewer sample values than is conventionally required. Consider the impulse response of a very wide-band system that has a two dominant medium-Q resonance's spaced widely in frequency at 10Hz and 10kHz. Conventionally, this type of wide-band impulse response must be linearly sampled at the higher rate set by the high-frequency content of the signal. The total time that the signal must be sampled however, is set by the much longer duration of the low-frequency ring down of the signal. This results in a unavoidably large number of sample points because the low-frequency content of the signal is highly over sampled.

If the effective sample rate is made higher at times near the start of the impulse and lower at times near the end of the impulse, the number of sample points can be vastly reduced. This is exactly what happens in log sampling [1].

With this type of filter, very-wide-band impulse responses can be accurately generated with very few sample points.

4. LOG-SAMPLED TYPE 2: INTERPOLATION IN FREQUENCY

4.1. Description

The logarithmically-sampled Type 2 filter allows direct specification of the real and imaginary parts of the filter's frequency spectrum in log frequency. As with the previous filter, the taps of this filter are spaced logarithmically in time. Each interpolated log sample in the filter's frequency response corresponds to a filter tap whose impulse response is a tone burst with a fixed number of cycles of varying duration, i.e. long at low frequencies and short at high frequencies. This time response is essentially the same as a constant-shape wavelet that is dilated or contracted in time. Fig. 5 illustrates this behavior.

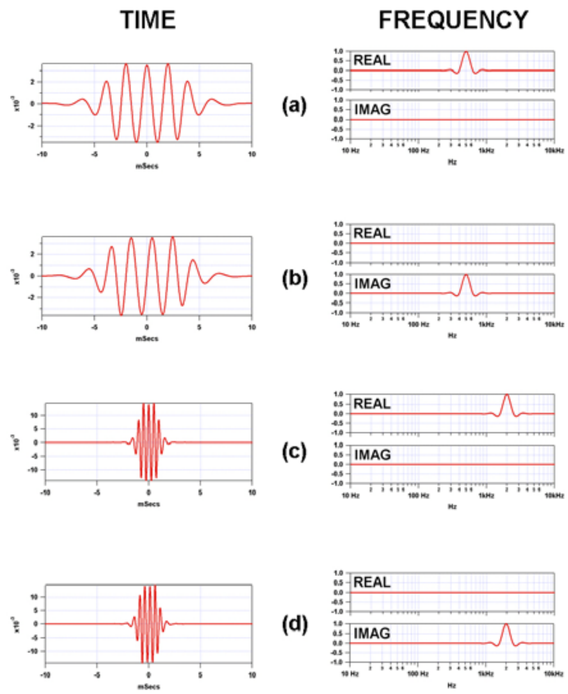

Fig. 5. Illustration of the time response of a cubic-spline interpolator in log frequency. Note *log* frequency scale of 10 Hz to 10 kHz (right), and *linear* time scale of –10 to +10 ms (left). The frequency sample rate is 10 points per decade, i.e. the interpolator steps in frequency every one-tenth decade (about one-third octave). (a) 500 Hz interpolator in real part of spectrum. (b) 500 Hz interpolator in imaginary part of spectrum. (c) 2,000 Hz interpolator in real part of spectrum. (b) 2,000 Hz interpolator in imaginary part of spectrum. Note that the time burst width decreases as frequency increases, but contains the same number of cycles, i.e. the waveform of the burst stays the same and dilates or contracts as frequency shifts.

With this technique, a 20 Hz to 20 kHz one-third-octave FIR equalizer can be implemented with only 31 log-spaced taps.

Fig. 6 shows the block diagram of the log-sampled type 2 interpolation-in-frequency convolver. Two versions of this convolver are shown. One has the start of the bursts aligned at zero (top), the other has the bursts aligned at their centers (bottom). The latter generates a correct convolution, while the former adds an extraneous phase rotation that has to be corrected.

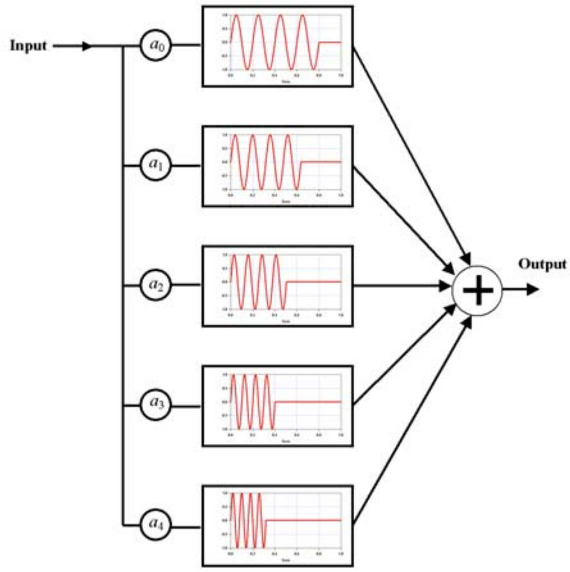

(a)

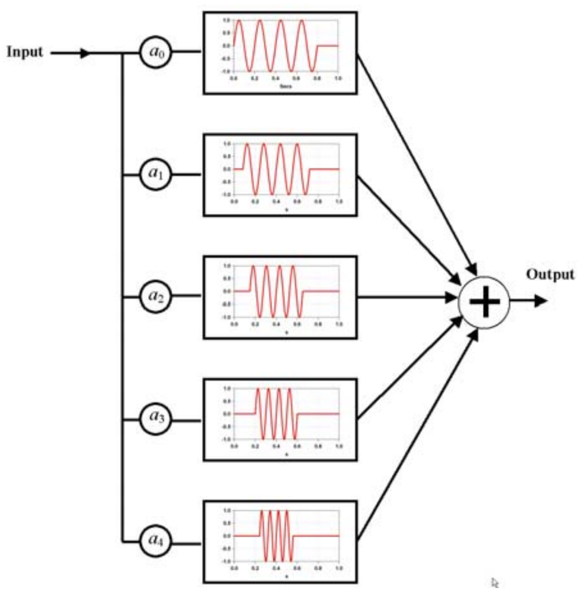

(b)

Fig. 6. Block diagram of a log-sampled type 2 interpolation-in-frequency FIR filter/convolver. Each tap's impulse response is a variable-width tone burst with a constant number of cycles in each burst. (a) All bursts aligned at start. (b) All bursts aligned at center. The frequency and phase of each burst depends on the location of the interpolator in the complex log frequency response. This form of convolver allows the filter's complex frequency response to be directly specified at a relatively small number of log-spaced points in frequency. Each channel is illustrated with a one second time-span convolver containing the variable-width tone burst of varying frequency. The illustrated bursts allow the log frequency response to be specified at one-tenth decade (essentially one-third octave) intervals with a sample ratio of $10^{0.1}$ = 1.2589. Five channels are illustrated with approximate log-spaced sample frequencies of 5, 6.3, 8, 10, and 12.5 Hz.

The convolver output is given by the following equation which assumes continuous time operation for simplification:

$$y(t) = \sum_{n=0}^{N-1} a_n x(t) \otimes \left[\Im^{-1} \left(H_{\text{interp}} \left(N_e \ln \left(\frac{f}{r^n f_{\min}} \right) \right) \right) \right] \quad (4)$$

where $x(t)$ is the input signal, $y(t)$ is the system output signal, N is the number of channels, the a_n are the amplitude coefficients of each channel, H_{interp} is the interpolation function, N_e is the sample density, r is the ratio between successive sample times ($r = e^{1/N_e}$), $f_{\min}$ is the sampling start frequency ($f_{\min}$ > 0), $\Im^{-1}$ is the inverse Fourier transform operator, and $\otimes$ is the convolution operator. Here the sampling stops (last frequency sample) at $f_{\max} = r^{N-1} f_{\min}$.

This type of convolver is conceptually similar to the previous linear convolver of Fig. 3, but each channel contains the time waveform of an inverse Fourier transformed log-warped interpolation function in frequency. Here however, the time waveforms are tone bursts of varying width and frequency but contain a constant number of cycles. This causes the bursts to be long at low frequencies and short at high frequencies.

As before, because of the complex frequency response, the burst frequency and phase changes at each tap position. Not indicated is Eq. 4, is an energy normalization term that keeps the energy of the tone bursts constant as their width decreases.

As an example, a 20-to-20 kHz pass-band one-third-octave equalizer can be created with only 31 log-spaced taps. Here the frequency step ratio is equal to the tenth root of ten ($f_{n+1}/f_n = 10^{0.1}$).

5. EFFICIENT IMPLEMENTATIONS

The log FIR convolvers of Figs. 4 and 6a may be more efficiently implemented with a discrete-time multi-rate structure as outlined in [1] for the type 1 interpolation-in-time convolver. These implementations are shown in Figs. 7 and 8 and the following description parallels that found in [1].

Each circle with an up arrow or down arrow indicates the operation of up sampling and down sampling respectively, by the fractional amount r. This operation may be implemented by IC chips such as the Analog Devices AD1890, AD1891, AD1893, or AD1896 asynchronous sample rate converters [8]. In these structures, each FIR filter block is identical, the changing sample rate automatically scales the interpolator's impulse response up and down in log time or the tone burst's frequency and width up and down in log frequency. The highest-sample-rate channel is at the top, while the lowest-sample-rate channel is at the bottom. Refer to the figure captions for more detailed information on these multi-rate convolvers. The weighting coefficients *a* allow direct one-to-one specification of the convolver's waveform in log time (Fig. 4) or the convolver's frequency response in log frequency (Fig. 6).

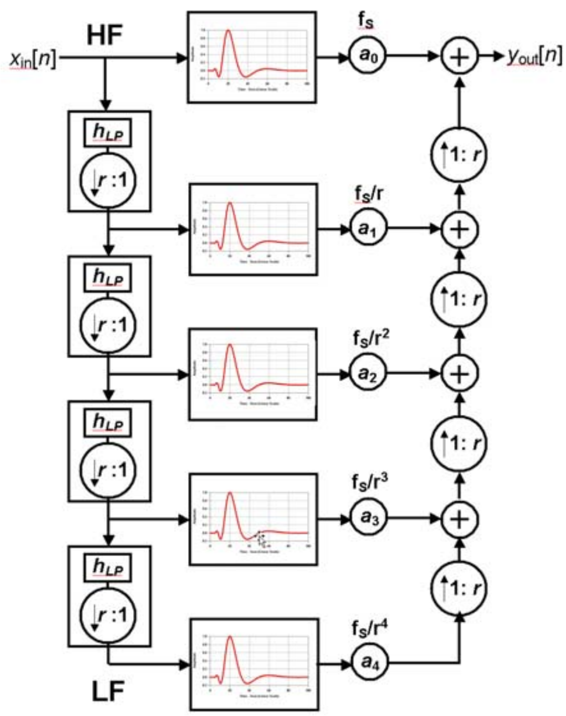

Fig. 7. Block diagram of a multi-rate implementation of the log-sampled type 1 interpolation-in-time FIR filter/convolver of Fig. 4. Each of the circular blocks containing an arrow provide fractional up and down sampling with rate ratio *r*. Each rectangular block on the left contains a low-pass filter followed by a down sampler. Each rectangular block in the center contains an FIR filter with a linear-sampled log-warped interpolation function. Each circle on the right contains an up sampler. Note that each channel of this convolver *is identical* except for the variable channel weights a_0 to a_4. The weights set the amplitude of each of the log-warped interpolator outputs and thus directly specify the convolvers waveform at a series of discrete points in log time. The required log time scaling of each channel is accomplished automatically by the different sample rates with the highest rate on top and the lowest rate on the bottom. Note that the structure can be extended indefinitely downward to accommodate arbitrarily low frequencies.

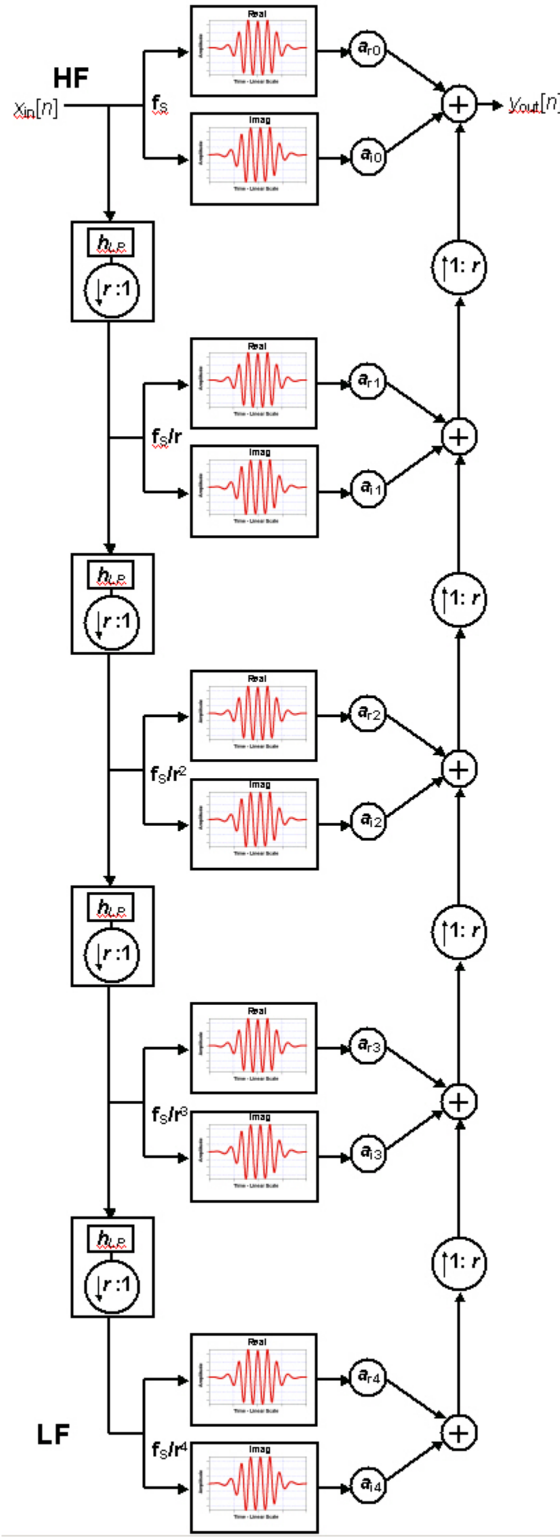

Fig. 8. Block diagram of a multi-rate implementation of the log-sampled type 2 interpolation-in-frequency FIR filter/convolver of Fig. 6a. Refer to the caption of Fig. 7 for an explanation of each blocks function. The rectangular blocks in the center contain an FIR filter with a linear-sampled tone burst. Each channel

however, includes two tone bursts each of which are associated with the real and imaginary interpolators in the frequency domain. The top waveform of each pair is a cosine burst associated with the real part of the spectrum at the burst's frequency, while the bottom waveform is a sine burst of the same frequency associated with the imaginary part of the spectrum. The bottom waveform is the Hilbert-transformed 90º-phase-shifted version of the top waveform. Note that each channel's waveform *is identical* to every other. The complex channel weights (a_{r0}, a_{i0}), (a_{r1}, a_{i1}) . . . (a_{rN}, a_{iN}) allow direct specification of the convolver's frequency response (real/imag or mag/phase) in a one-to-one correspondence between the burst's center frequency and the corresponding point in log frequency. As with Fig. 7, the structure can be extended indefinitely downward to accommodate arbitrarily low frequencies.

6. CONCLUSIONS

Traditional FIR filters designed to operate at the typically-high sample rates of traditional hardware, often require very-long filters to properly specify waveforms and frequency responses for wide-band signals. A wide-band signal in this context is defined as a signal that contains significant energy in frequencies that span a bandwidth greater than about one octave. Audio signals definitely fit this definition because of their very-wide three-decade 20 Hz to 20 kHz bandwidth. These long filters often provide an excessive amount of time or frequency resolution and exhibit high design complexity because a very-long FIR filter represents a system with an extremely-high number of degrees of freedom.

To get around these problems, IIR filters are often used for these situations, but these filters may exhibit problems at low frequencies because of the requirement for high precision arithmetic and the effects of round off errors.

This paper has presented an alternate means to get around these problems in a class of FIR/convolution filters based on interpolation. Four types of filter/convolvers were described: 1) linear sampled type 1: interpolation in time, 2) linear sampled type 2: interpolation in frequency, 3) log sampled type 1: interpolation in time, and 4) log sampled type 2: interpolation in frequency.

Linear interpolation allows the effective sample rate of the filters to be reduced when applied to the time domain or allows significantly fewer samples to specify a smoothly-varying frequency response when applied to the frequency domain.

Log interpolation, a relatively new concept, allows the thrifty specification of the impulse response of very-wide-band systems. In log-sampling or interpolation in time, the sample rate is high at the start of the sampling and then falls inversely with time reaching low values at long times. This is found to be a good match to systems composed of resonators of roughly the same Q, that cover a very-broad range of resonant frequencies. Log interpolation or sampling in frequency allows very sparse specification of frequency responses when applied to responses that are often viewed on a log frequency scale.

The convolvers described here take advantage of linear and log interpolation to greatly reduce the number of data points required to specify time waveforms for certain signals and to specify the frequency response of systems whose response slowly varies in linear frequency or that exhibit constant-Q behavior over a wide frequency range in log frequency. The convolver's design process is considerable simplified, because all the convolvers provide a one-to-one correspondence between the filter tap weights and the resultant time response or frequency response in either the linear or log domains.

An alternate definition of the sampling and reconstruction process was reviewed which was based on the orthogonal characteristics of the sinc function. This process essentially combines the filtering and sampling of the conventional method in one operation. The alternate sampling method allows straightforward calculation of samples in either linear or log time or in linear or log frequency.

Although some of the described convolvers are very inefficient when compared to standard FIR filters, the interpolation methods presented may provide alternate practical ways to describe and implement some systems. The described convolvers can have far fewer samples or FIR taps that conventional FIR filters, but do require much more complex processing at each tap or convolver channel. A efficient implementation method was described for the log-sampled convolvers that relies on multiple fractional-rate asynchronous sample rate converters.

7. REFERENCES

[1] D. B. Keele, Jr, "Log Sampling in Time and Frequency: Preliminary Theory and Application," Presented at the 97th convention of the Audio Eng. Soc., preprint no. 3935 (Nov. 1994).

[2] Hatziantoniou, D. Panagiotis, and J. N. Mourjopoulos, "Generalized Fractional-Octave Smoothing of Audio and Acoustic Responses," J. Audio Eng. Soc., vol. 48, no. 4, pp. 259-280 (April 2000).

[3] M. R. Bai and P. Zeung, "Fast Convolution Technique Using a Nonuniform Sampling Scheme: Algorithm and Applications in Audio Signal Processing" J. Audio Eng. Soc., vol. 52, no. 5 (May 2004).

[4] J. O. Smith and P. Gossett, "A Flexible Sampling-Rate Conversion Method," Proc. IEEE Int. Conf. On Acoustics, Speech, and Signal Processing, San Diego, (March 1984). Also see (http://ccrma.stanford.edu/~jos/resample/resample.html)

[5] F. J. Harris, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," Proc. IEEE, vol. 66, no. 1, pp. 51-83 (Jan. 1978).

[6] P. G. Craven, "Anti-Alias Filters and System Transient Response at High Sample Rates," J. Audio Eng. Soc., vol. 52, no. 3 (Mar. 2004).

[7] M. Unser, A. Aldroubi, and M. Eden, "Fast B-spline Transforms for Continuous Image Representation and Interpolation," IEEE Pattern Anal. Machine Intell., vol. 13, pp. 277-285 (1991).

[8] http://www.analog.com/

8. APPENDIX 1: AN ALTERNATE VIEW OF SAMPLING

An alternate more rigorous definition of the sampling and reconstruction process can be formulated using the orthogonal characteristics of the sinc function:

$$\int_{-\infty}^{\infty} \text{sinc}(x-j)\text{sinc}(x-k)dt = \begin{cases} 1 & \text{for } j = k \\ 0 & \text{for } j \neq k \end{cases} \tag{5}$$

where

$$\text{sinc}(x) = \sin(\pi x)/\pi x$$
$$j,k \quad = \text{integers}$$

These equations state that if you multiply two aligned sinc functions and then integrate you get unity, but zero if the two are unit shifted and not aligned.

In 1994 I outlined an alternate definition of sampling and reconstruction process which takes advantage of the above orthoganality of the sinc function in an AES paper I gave to the 97th convention [1, Section 2.1].1 The following description is based on that paper.

Conventional sampling theory assumes that the continuous-time input signal is first band limited so that the signal contains no frequency components at or above one-half the sampling frequency f_S. This band limited continuous-time signal is in turn converted to sequence of samples by passing the signal first through an impulse-train modulator operating at f_S and then converting from the series of weighted impulses to a series of numbers representing the sample values at each instance of time.

The alternate view of the sampling operation is to view it as the decomposition of the input continuous-time signal into a series of shifted and weighted sinc functions. The amount of shift is equal to an integral multiple of the sample time period (= $\pm nT = \pm n/f_S$, where n is an integer).

The decomposition appears as:

$$x[n] = \int_{-\infty}^{\infty} f(t)\sin c(t-nT)dt \tag{6}$$

where $x[n]$ is a sequence of numbers that represents the sample amplitudes at each instance of time, n is the sample count integer, $f(t)$ is the continuous time function to be sampled, T is the sampling period ($1/f_S$), and f_S is the sampling frequency in Hz.

This alternate view of the sampling can be thought of as an operation that combines the low-pass filtering and sampling of the conventional method in one operation.

The conventional reconstruction process, where the sample numbers are converted to a sequence of impulses at the sample rate and then passed through an ideal reconstruction filter, is recognized as the inverse process of the previous sampling operation of Eq. (6). This is shown in the following:

$$x_r(t) = \sum_{n=-\infty}^{+\infty} x[n]\,\text{sinc}\left[(t-nT)/T\right] \tag{7}$$

where

$x_r(t) =$ reconstructed continuous-time signal

$x[n] =$ sequence of numbers representing amplitudes at each instance of time

$n =$ sample count, integer

$T =$ sample period ($=1/f_S$)

$f_S =$ sample rate, Hz

The sampling process of Eq. (6) decomposes the continuous-time input signal into a discrete-time signal which is a series of weights or samples of the shifted sinc functions. The reconstruction process of Eq. (7) multiplies the weights or samples by each of the individual shifted sinc functions and sums each product to reconstruct the continuous time signal. The sinc function acts as an interpolation function that serves to fill in the data between the samples.

The idea for this alternate view of sampling came to me after reading several texts on wavelets. This alternate sampling viewpoint is much akin to the wavelet decomposition process where the wavelet is a sinc function. The problem with the sinc function is that it is infinite in extent. In wavelets terms it is said to not have "compact support." The whole field of dyadic (base 2) wavelet theory is based on discovering wavelets that do have compact support but are orthogonal in the sense of Eq. 2. Many have been discovered since 1988 that many believed did not exist.

The situation is similar in sampling theory where the quest is to find a reasonable sampling or interpolation function that is finite in extent but provides the proper roll off in the opposite domain. A windowed sinc function is appropriate for sampling because it does have compact support. A Hann-windowed sinc function centered at $x = 0$ appears as

$$\text{Hannsinc}(x,W) = \begin{cases} \left[1+\cos\left(\dfrac{2\pi x}{W}\right)\right]\text{sinc}(x) & \text{for } -\dfrac{W}{2} \le x \le +\dfrac{W}{2} \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where W is the width or support range of the Hannsinc and is an even integer in the range of 2 and higher. Higher values of W extend the width of the Hannsinc and increase the sharpness of the edges of the rectangular function in the opposite domain.

Other windows such as the Keiser-Bessel are appropriate as well [4][5, p. 73]. A truncated cubic-spline interpolator is also appropriate as an interpolator [1, Appendix 2][6].

9. APPENDIX 2: LOG SAMPLING

Log sampling in time is a non-uniform sampling technique where the sample rate is high initially and then falls inversely with time. The sample times form a geometric sequence with a constant ratio ($r > 1$) from each sample time to the next, i.e 0.1, 0.2, 0.4, 0.8, etc for $r = 2$. As with the previous appendix, the following explanation is based on [1].

The technique is optimum for sampling the impulse response of real-world systems that are often composed of roughly constant-Q resonators that have short impulse responses at high-frequencies and long impulses at low frequencies. This is roughly the case for systems such as loudspeakers and the reverberant acoustic decay of enclosures.

The log sampling technique can also be applied to the frequency domain where the frequency response sampling has a high rate at low frequencies and a low rate at high frequencies. This is entirely equivalent to first applying a constant-Q smoothing to a frequency response and then sampling the frequency response at constant percentage octave- or third-octave increments.

Log sampling in time applies a warping function to time that maps linear time to log time using the following warping function based on natural logarithms:

$$\tau = N_e \ln(t / t_c) \quad (9)$$

or base-10 logarithms:

$$\tau = N_{10} \log(t / t_c) \quad (10)$$

where τ is the log-time variable, t is the linear-time variable, N_e and N_{10} are constants that specify a zoom factor or sample density (in points per e [= 2.7183] or points per decade [= 10]), and t_c is the center time, i.e. the linear time value that corresponds to zero in log time.

The sample-density constants N_e and N_{10} indicate how many samples will occur in a specific percentage increment of log time. It essentially specifies the amount of shift that the interpolator must be shifted in log time from one sample to the next. The ratio r between the sample times of successive samples is given by

$$r = e^{1/N_e} = 10^{1/N_{10}} \quad (11)$$

For example: a N_{10} value of 10 points per decade indicates a one-tenth-decade shift ($r = 10^{1/N_{10}} = 10^{1/10} = 10^{0.1} \approx 1.26$) or about a one-third-octave shift in time between the time of one sample to the next.

This transformations of Eqs. (9) and (10) maps the positive linear time axis $(0 > t > +\infty)$ onto the whole log τ axis $(-\infty > \tau > +\infty)$, with times greater than t_c mapped to the positive τ axis and times less than t_c mapped to the negative τ axis. This transformation effectively expands the time axis for short times, and compresses the time axis for long times.

When the warping functions of Eqs. (9) and (10) are converted to discrete time, the effective sample rate can be shown to fall inversely with time [1, Eq. (14)] as follows:

$$f_S(t) = \frac{1}{(r-1)t} \tag{12}$$

where r is the ratio between the time of successive samples. In log time, the sampling process extends from $t_{\min}$ to $t_{\max}$. The sampling rate is maximum at $t_{\min}$ and then falls inversely with time reaching a minimum at $t_{\max}$.

With the mapping functions of Eqs. (9) and (10), the Hannsinc can be warped to operate in log time (using the base-10 logarithm version):

$$\text{LogHannsinc}(t, t_c, N_{10}, W) = \text{Hannsinc}\left[N_{10} \log\left(\frac{t}{t_c}\right), W\right] \tag{13}$$

where t_c is the center time of the of the interpolator in log time, N_{10} is the sample density, and W is the width of the LogHannsinc in log time in terms of the number of zero crossings (W = 2, 4, 6,…).

The LogHannsinc function is unity at $t = t_c$ and has zero crossings at every percentage multiple above and below t_c, i.e., $\ldots \frac{t_c}{r^3}, \frac{t_c}{r^2}, \frac{t_c}{r}, rt_c, r^2 t_c, r^3 t_c, \ldots$ etc., and only exists for $t > 0$ (by definition). The LogHannsinc acts as an interpolation function which is specifically designed to interpolate a sequence of samples whose sample times are arranged in a geometric sequence with r the ratio between successive sample times.

Figure 9 shows the waveform of a base-two (N_{10} = 1/log(2) ≈ 3.32 points per decade) LogHannsinc in time plotted on both linear and log time scales. When plotted on a linear-time scale, the waveform is compressed on the left and expanded on the right of the peak. However, when plotted on a log time scale the waveform is completely symmetrical and is designed to interpolate a sequence of samples at log time values forming a base-two (r = 2) geometric sequence.

With a substitution of frequency for time, the LogHannsinc can operate in log frequency as well.

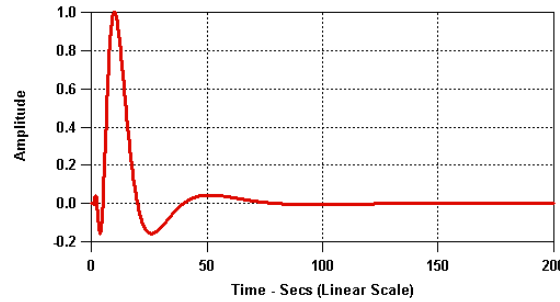

(a)

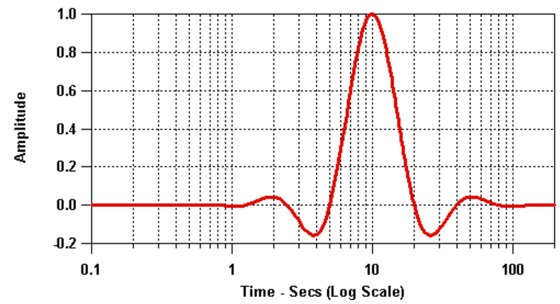

(b)

Fig. 9. Plot of a base-two LogHannsinc interpolator in time. (a) Linear time scale of 0 to 200 s. (b) Log time scale of 0.1 to 200 s. The LogHannsinc is centered at 10 s and has a width of 8, i.e. its support range extends from 0.625 to 160 s. Zeros of the function occur at 1.25, 2.5, 5, 20, 40, and 80 s, and has a peak of unity at 10 s.