ON MINIMIZING THE LOOK-UP TABLE SIZE IN QUASI-BANDLIMITED CLASSICAL WAVEFORM OSCILLATORS

Jussi Pekonen†, Juhan Nam‡, Julius O. Smith‡, Jonathan S. Abel‡ and Vesa Välimäki†

†Dept. of Signal Processing and Acoustics
Aalto University
Espoo, Finland
Jussi.Pekonen@tkk.fi
Vesa.Valimaki@tkk.fi

‡Center for Computer Research in Music and Acoustics
Stanford University
Stanford, California 94305, USA
juhan@ccrma.stanford.edu
jos@ccrma.stanford.edu
abel@ccrma.stanford.edu

ABSTRACT

In quasi-bandlimited classical waveform oscillators, the aliasing distortion present in a trivially sampled waveform can be reduced in the digital domain by applying a tabulated correction function. This paper presents an approach that applies the correction function in the differentiated domain by synthesizing a bandlimited impulse train (BLIT) that is integrated to obtain the desired bandlimited waveform. The ideal correction function of the BLIT method is infinitely long and in practice needs to be windowed. In order to obtain a good alias-reduction performance, long tables are typically required. It is shown that when a short look-up table is used, a windowed ideal correction function does not provide the best performance in terms of minimizing aliasing audibility. Instead, audibly improved alias-reduction performance can be obtained using a look-up table that has a parametric control over the low-order generations of aliasing. Some practical parametric look-up table designs are discussed in this paper, and their use and alias-reduction performance are exemplified. The look-up table designs discussed in this paper providing the best alias-reduction performance are parametric window functions and least-squares optimized multiband FIR filter designs.

1. INTRODUCTION

In "subtractive" sound synthesis, the spectrum of a spectrally rich source signal is shaped with a time-varying resonant filter. This sound synthesis approach was a popular choice in the analog synthesizers of the 1960s and 1970s. Traditionally, the source signals used in subtractive synthesis included sawtooth, rectangular pulse, and triangular pulse waves [1, 2]. However, simply sampling these classical geometric waveforms suffers from harsh aliasing caused by the discontinuities in the waveform or its derivative. This aliasing problem is illustrated in Fig. 1 where a sawtooth wave having fundamental frequency $f_0 = 2.637$ kHz (note E7) is shown together with the spectrum of its trivially sampled digital representation for the sampling frequency $f_s = 44.1$ kHz. The desired harmonics below the Nyquist limit are indicated with crosses. The examples presented in this paper use the same frequencies.

Several approaches to reduce or remove the aliasing distortion have been suggested. Some of these algorithms are truly bandlimited as they synthesize only a fixed number of harmonics at each time instant. These algorithms include additive synthesis [3] and its related approaches "discrete summation formulae" [4, 5] and inverse-fast-Fourier-transform synthesis [6], wavetable synthesis [7], and triggering of a feedback delay loop [8].

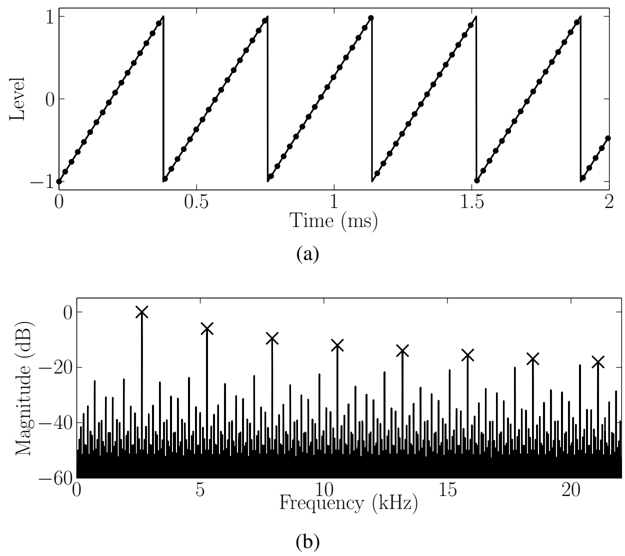

Figure 1: (a) *Waveform of a sawtooth waveform having fundamental frequency f_0 = 2.637 kHz (note E7) and* (b) *the spectrum of its trivially sampled digital representation. A sampling frequency f_s = 44.1 kHz was used. The sampled data is marked with dots in* (a). *The desired harmonics, i.e., the components that a bandlimited signal would contain, are indicated with crosses in* (b).

In another group of algorithms, some aliasing is allowed at high frequencies where human hearing is less sensitive. These algorithms apply a pre-sampling lowpass filter either to the waveform [9] or its derivative (an impulse train) [10, 11]. Approximately bandlimited impulses can also be generated using modified FM synthesis [12] or fractional delay filters [13, 14].

Other algorithms produce aliasing over the whole audio band, but the aliasing is suppressed relative to the trivially sampled case. Examples include oversampling [15, 16], differentiated polynomial waveforms [17–20], filtering of a distorted sinusoid [21], and amplitude and phase distortion synthesis approaches [22–28]. Some

This work has been partly funded by the Academy of Finland, project number 122815.

of the distortion synthesis techniques can be bandlimited but generally they are not. Alternatively, the aliasing of the trivially sampled waveform can be reduced in some cases by applying a post-processing filter in the digital domain [29]. For example, exposed aliasing terms can be notched out or highpass-filtered out when they lie below the fundamental frequency. Such post-processing filters can also be applied to the output of the various algorithms mentioned above.

In this paper, the group of alias-reduction algorithms that perform lowpass filtering prior to sampling is investigated in more detail. Section 2 reviews the operation principles of these quasi-bandlimited waveform-synthesis algorithms. In Section 3, minimization of the look-up table size utilized in quasi-bandlimited oscillator algorithms is discussed, and Section 4 presents a set of practical look-up table design approaches. Finally, Section 5 concludes the paper and points out some directions for further work.

2. QUASI-BANDLIMITED OSCILLATOR ALGORITHMS

In quasi-bandlimited waveform synthesis, the waveform to be synthesized can be understood to be lowpass filtered prior to sampling. In practice, these algorithms apply a correction function in the digital domain, resulting effectively in the same result. In the ideal case, the correction function can be expressed in closed form, and its derivation is reviewed next.

One quasi-bandlimited approach operates on the waveform derivative. For instance, the continuous-time sawtooth waveform of Fig. 1(a) can be expressed in as

$$s(t) = 2tf_0 - 2\sum_{k=-\infty}^{\infty} u(t - kT_0), \tag{1}$$

where f_0 is the fundamental frequency in Hertz, $T_0 = 1/f_0$ is the oscillation period in seconds, and $u(x)$ is the Heaviside unit step function,

$$u(x) = \begin{cases} 0, & \text{for } x < 0, \\ 0.5, & \text{for } x = 0, \text{and} \\ 1, & \text{for } x > 0. \end{cases} \tag{2}$$

The differentiated sawtooth waveform is given by

$$\frac{\mathrm{d}}{\mathrm{d}t}s(t) = 2f_0 - 2\sum_{k=-\infty}^{\infty} \delta(t - kT_0), \tag{3}$$

where $\delta(x)$ is the Dirac delta function, the derivative of the unit step function $u(x)$, which is zero for $x \neq 0$ and has the property

$$\int_{-\infty}^{\infty} \delta(x)\mathrm{d}x = \int_{0-}^{0+} \delta(x)\mathrm{d}x = 1. \tag{4}$$

By bandlimiting the impulse train in the differentiated domain and integrating the resulting bandlimited impulse train (BLIT) that contains a DC offset $2f_0$ (see Eq. (3)), an approximately bandlimited sawtooth waveform is generated [10,11]. In BLIT, every Dirac delta function in the waveform derivative of Eq. (3) is replaced with the impulse response of the applied lowpass filter [10,11]. In the ideal case, the lowpass filter impulse response is given by

$$h_{\mathrm{id}}(t) = \mathrm{sinc}(2f_{\mathrm{c}}t), \tag{5}$$

where $\mathrm{sinc}(x) = \sin(\pi x)/(\pi x)$ and f_{c} is the cutoff frequency of the lowpass filter.

It should be noted that since the sinc function is infinitely long, an accurate BLIT synthesizer should sum infinitely many sinc-function values at each sampling instant. In order to have an implementable realization, the ideal correction function must be truncated somehow. In practice, an appropriate portion of the sinc function is windowed in order to avoid discontinuities at the truncation boundaries [10]. This means that the ideal lowpass filter response is replaced with a function

$$h(t) = w(t)\mathrm{sinc}(2f_{\mathrm{c}}t), \tag{6}$$

where $w(t)$ is the applied continuous-time window function that is non-zero on a finite interval. However, inline evaluation of this function can be computationally expensive, so it is typically pre-computed and stored to a look-up table that is triggered for each discontinuity [30].

When the impulses of Eq. (3) are replaced with $h(t)$ and the resulting signal is sampled, another issue can be observed. The sampled and time-shifted $h(t)$ yields

$$h(n - D) = w(n - D)\mathrm{sinc}(n - D), \tag{7}$$

where n is the sample index and D is the delay, which is an integer multiple of the oscillation period in samples. Now, since the oscillation period is in general not an integer, the mid-point of the look-up table is not located on a specific sampling instant. This means that the look-up table needs to be oversampled or interpolated to obtain proper positioning of the look-up table for each discontinuity. Both oversampling and interpolation can be used to further improve the accuracy. If a high enough oversampling factor is used, linear interpolation suffices. In the examples of this paper, linear interpolation is incorporated.

From the discussion above, the look-up table used in the BLIT algorithm can be summarized to be a NMth-order windowed sinc function FIR filter having the normalized cutoff frequency

$$\omega_{\mathrm{c}} = \frac{2\pi f_{\mathrm{c}}}{Mf_{\mathrm{s}}}, \tag{8}$$

with $M \in \mathbb{N}$ denoting the oversampling factor and $N \in \mathbb{N}$ denoting the number of samples the algorithm modifies in a time interval about each waveform discontinuity.

In the BLIT algorithm, the synthesized bandlimited impulse train needs to be integrated in order to obtain the actual waveform. However, when the BLIT algorithm is implemented with finite accuracy, numerical integration with an accumulator filter may lead to numerical problems. Brandt proposed in [9] to use a second-order leaky integrator that has zero DC gain,

$$H_{\mathrm{int},2}(z) = \frac{\pi(1+c)}{2}\frac{1 - z^{-1}}{(1 - cz^{-1})(1 - cz^{-1})}, \tag{9}$$

where c sets the leakage. With leakage factor $c = 0.9992$ the filter has 0.5 dB attenuation at 20 Hz when the sample rate is 44.1 kHz. In the examples of this paper, this second-order leaky integrator with leakage factor $c = 0.9992$ is used.

Furthermore, Brandt suggested that the integration could be performed before the sampling [9]. The approach Brandt suggested uses an accumulated windowed sinc function as an approximation of the bandlimited step function (BLEP). The BLEP table is triggered at each discontinuity and output, and when the table read is finished, the BLEP synthesizer outputs a unit-amplitude constant. However, if the unit step function is subtracted from the table, the resulting BLEP residual can be utilized as a correction function that is added onto the trivial, non-bandlimited waveform at each discontinuity [30, 31].

3. MINIMIZING THE LOOK-UP TABLE SIZE

It is advantageous to use as short a look-up table size as possible in order to minimize memory consumption. This also helps to minimize overlapping reads of the look-up table at high fundamental frequencies [22]. However, the alias-reduction performance typically improves as the number of modified samples increases. Furthermore, as the required aliasing reduction is different for different fundamental frequencies, the choice of optimal look-up table parameters is not a trivial task.

When a short table size is used, the use of a look-up table that is *not* the windowed sinc function may provide a better aliasing-reduction performance, as illustrated in Fig. 2 for look-up table parameters $N = 4$, $M = 8$, and $f_\mathrm{c} = 0.5 f_\mathrm{s}$. It can be seen in Fig. 2(a) that the spectrum of a sawtooth wave synthesized with the BLIT algorithm using a Hann-windowed sinc function contains considerable aliasing between the harmonics and below f_0. When a plain Hann window is used instead of the windowed sinc function by the BLIT algorithm, the aliasing is reduced, see Fig. 2(b). However, this improvement in aliasing reduction comes at the price of attenuation of the higher components. Both phenomena can be explained with the help of the magnitude responses of these look-up tables given in Fig. 2(c). It can be seen that the Hann window suppresses the first-order generation of aliasing (frequencies from $0.5 f_\mathrm{s}$ to f_s) more than the windowed sinc function, but it also suppresses the desired passband (frequencies below $0.5 f_\mathrm{s}$), whereas the windowed sinc function does not suppress signal components in the passband.

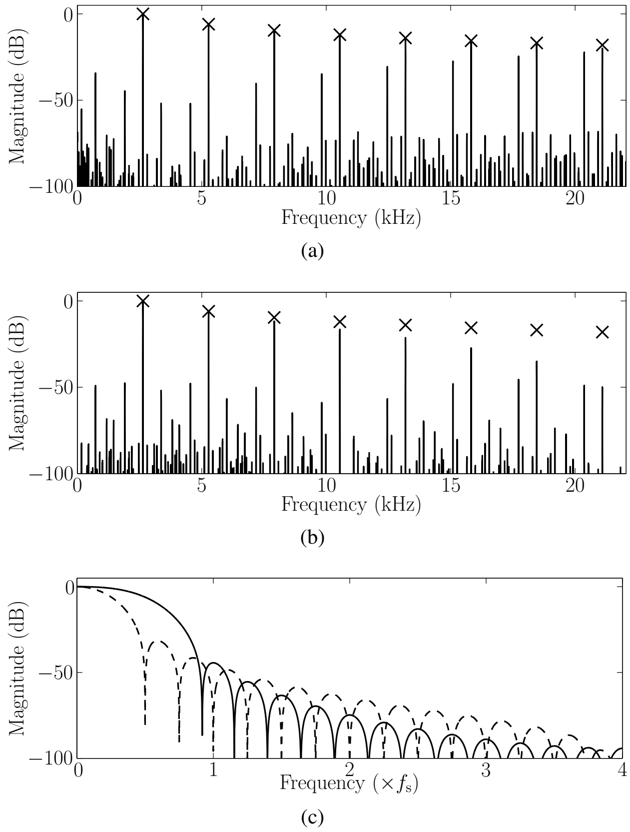

Figure 2: *Spectra of the sawtooth waves synthesized with the BLIT algorithm using* (a) *a Hann-windowed* sinc *function and* (b) *a plain Hann window. The amplitude drop of the higher components with the plain Hann window is clearly shown in* (b)*. The desired amplitudes of the non-aliased components are marked with crosses. The magnitude responses of the Hann-windowed* sinc *(solid line) and the plain Hann window (dashed line) look-up tables are shown in* (c)*.*

3.1. Perceptually optimal look-up tables

The example above indicates that the look-up table providing the best aliasing reduction while being as short as possible is likely not the windowed sinc function, because the difference in aliasing audibility would appear to be larger than the difference in the audibility of attenuated harmonic components near the Nyquist limit. The amount of aliasing reduction depends on how fast the look-up table's magnitude spectrum rolls off above the cutoff frequency: Less aliasing requires more attenuation of high-frequency harmonics. Accordingly, one definition of a perceptually optimized design is a look-up table that has a minimized length subject to the constraint that both aliasing and spectral-reshaping (rolloff in the example above) are inaudible.

In general terms, the search for the perceptually optimal look-up table $\bar{g}(k)$, where k is the table index, for a given fundamental frequency f_0, can be understood as an optimization problem that minimizes the size of a look-up table design $g(k)$, expressed as

$$\hat{L}(g(k), f_0) = \min NM \tag{10}$$

for $N, M \in \mathbb{N}$ subject to constraints

$$E_A = L_A(g(k), f_0, f) - \theta_A(g(k), f_0, f) \leq 0, \tag{11}$$

$$E_{H,k} = \theta_H(kf_0) - L_H(g(k), f_0, k) \leq 0, \tag{12}$$

where $L_A(g(k), f_0, f)$ is the level of the aliased component at frequency f, $\theta_A(g(k), f_0, f)$ is the threshold of audibility of aliasing at frequency f, $\theta_H(kf_0)$ is the maximum allowable amplitude drop at the frequency of the kth harmonic component, and $L_H(g(k), f_0, k)$ is the level of the kth harmonic component. The constraints of Eqs. (11) and (12) can be included in the objective function of Eq. (10) via a penalty function transform [32], yielding

$$\hat{L}(g(k), f_0) = \min NM + \mu_0 \max\{E_A, 0\} + \sum_{k=1}^{\lfloor \frac{f_\mathrm{s}}{2f_0} \rfloor} \mu_k \max\{E_{H,k}, 0\}, \tag{13}$$

where μ_0 and μ_k are large numbers, e.g., 10^5 or 10^6, that force the constraints to be satisfied.

It should be noted that L_A, θ_A, and L_H depend on the fundamental frequency f_0 and look-up table design $g(k)$. Therefore, this optimization problem provides the optimized table size only for a specific fundamental frequency using the given look-up table design. The optimal table size $\bar{L}$ for a look-up table design $g(k)$ and for a range of interesting fundamental frequencies F_0 is obtained as the maximum of these fundamental-frequency-specific table sizes, i.e.,

$$\bar{L}(g(k)) = \max_{f_0 \in F_0} \hat{L}(g(k), f_0). \tag{14}$$

In order to have a feasible solution, the look-up table design $g(k)$

having the table size $\bar{L}(g(k))$ should satisfy the constraint of Eqs. (11) and (12) for all $f_0 \in F_0$. The general, perceptually optimal look-up table $\bar{g}(k)$ is the look-up table design $g(k)$ that has the smallest table size $\bar{L}(g(k))$ in F_0. However, finding a feasible solution to this optimization problem is, in practice, a very complex task due to the highly nonlinear and implicit formulation. In addition, it is not guaranteed that the problem even has a feasible solution. A simpler alternative is a parametric approach that provides a control over the stopband attenuation that, in the ideal case, would be independent of the filter order.

4. PRACTICAL LOOK-UP TABLE DESIGNS

4.1. Parametric window functions

As pointed out above, the aliasing can be reduced by utilizing a window function as the look-up table used by the BLIT algorithm. However, typical window functions such as the Hann window are non-parametric, and their stopband attenuation—the amount of aliasing suppression—is constant. Nevertheless, there exist window functions that offer parametric control over the minimum stopband attenuation α_s. Two commonly known examples are the Kaiser window and the Dolph-Chebyshev window.

The Kaiser window of length $L = NM + 1$ is given by [33]

$$w_K(k) = \frac{I_0\left(\beta\sqrt{1 - k^2/(L-1)^2}\right)}{I_0(\beta)} \tag{15}$$

for $k = -NM/2, \ldots, NM/2$, where $I_0(x)$ is the zeroth-order Bessel function given by

$$I_0(x) = 1 + \sum_{k=1}^{\infty} \frac{x^{2k}}{2^{2k}(k!)^2}, \tag{16}$$

and β is the parameter that controls the minimum stopband attenuation α_s:

$$\beta = \begin{cases} 0.1102(\alpha_s - 8.7), & \text{for } \alpha_s > 50, \\ 0.5842(\alpha_s - 21)^{0.4} + 0.07886(\alpha_s - 21), & \\ & \text{for } 21 \le \alpha_s < 50, \\ 0, & \text{for } \alpha_s < 21 \end{cases} \tag{17}$$

The Dolph-Chebyshev window of length $L = NM + 1$ is expressed as [34]

$$w_C(k) = \frac{1}{L}\left[\gamma + 2\sum_{m=1}^{(L-1)/2} T_{L-1}\left(\tau\cos\left(\frac{m\pi}{L}\right)\right) \times \cos\left(\frac{2km\pi}{L}\right)\right] \tag{18}$$

for $k = -NM/2, \ldots, NM/2$, where

$$\gamma = 10^{\alpha_s/20}, \tag{19}$$

$$\tau = \cosh\left(\frac{\cosh^{-1}(\gamma)}{L-1}\right), \tag{20}$$

and $T_l(x)$ is the lth-order Chebyshev polynomial defined as

$$T_l(x) = \begin{cases} \cos(l\cos^{-1}(x)), & \text{for } |x| \le 1, \\ \cosh(l\cosh^{-1}(x)), & \text{for } |x| > 1. \end{cases} \tag{21}$$

The effect of the window and look-up table parameters on the resulting spectrum of a sawtooth waveform synthesized with the BLIT algorithm using the Kaiser window is shown in Fig. 3. With a stopband attenuation $\alpha_s = 110$ dB and look-up table parameters $N = 4$ and $M = 8$, the resulting spectrum contains less aliasing than the spectrum obtained with the Hann-windowed sinc function or with the plain Hann window, especially at low frequencies (compare Fig. 3(a) with Figs. 2(a) and 2(b)). In addition, the amplitude drop of the higher desired components is lower than with the Hann window (see Fig. 2(b)).

With the parametric window functions, the amplitude drop of the higher components can further be compensated by increasing the minimum stopband attenuation α_s (see Fig. 3(b)). However, this compensation would also amplify the first-order generation of aliasing as the stopband attenuation is traded off for the width of the transition band. The larger the stopband attenuation, the broader the transition band, and more high-power aliasing will be passed without attenuation. Increasing N further suppresses the aliasing, but it also results in a larger amplitude drop of the higher desired components (see Fig. 3(c)). Similar observations can be made with respect to the Dolph-Chebyshev window—see Fig. 4.

The amplitude drop of the higher desired components can also be corrected with a post-processing equalizing filter that amplifies the high frequencies. Now, as the aliasing occurs in the synthesis stage and the equalizing filter boosts only the high frequencies, the level of aliasing at low frequencies remains approximately unchanged. However, in subtractive sound synthesis these classical waveforms are typically filtered with a lowpass filter, meaning that for small amplitude drops the equalizing filter is not a necessary operation. Moreover, the equalizing filter can be designed to amplify the higher components up to, e.g., 15 kHz above which human hearing is very insensitive [35].

If high-frequency compensation is desired, it could be obtained with a first-order IIR filter, given by

$$H_c(z) = \frac{1-p}{1-v}\frac{1 - vz^{-1}}{1 - pz^{-1}}, \tag{22}$$

where p and v are the locations of the filter pole and zero, respectively. The filter has unity gain at dc and a controllable gain for other frequencies. However, it should be noted that the gain as a function of frequency can be set with two reference points, e.g., 1 kHz and 15 kHz, meaning that the amplitude drop of the higher harmonics cannot be compensated *accurately*, but the compensation can be set within couple of decibels compared to the desired amplitudes. Alse, at such high frequencies there are normally many harmonics per critical band of human hearing, and the needed granularity of equalization accuracy is on the order of a critical bandwidth.

Using the approach described above, the amplitude drop of the table constructed from a Kaiser window having parameters $N = 4$, $M = 8$, and $\alpha_s = 110$ dB can be compensated to be within one decibel with coefficients $p = -0.2864$ and $v = 4.5159$. For a Dolph-Chebyshev window having the same parameters, the compensation to within one decibel is obtained with coefficients $p = -0.3092$ and $v = 4.6424$.

4.2. FIR filter design via optimization

Alternatively, the look-up table can be designed by approximating a target amplitude response $D(\omega)$, where ω is the normalized frequency variable, so that the resulting look-up table minimizes a

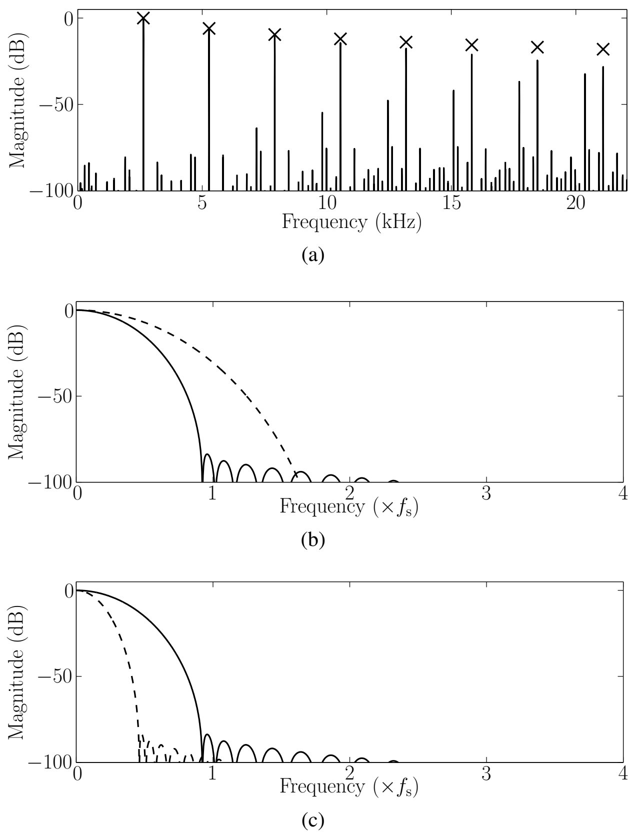

Figure 3: (a) *Spectrum of the sawtooth waveform synthesized with the BLIT algorithm using a Kaiser window with parameters* $N = 4$*,* $M = 8$*, and* $\alpha_s = 110$ *dB. The effects of the stopband attenuation and number of samples to be corrected on the aliasing reduction are illustrated in* (b) *and* (c)*, respectively, using table parameters* $N = 4$*,* $M = 8$ *and* $\alpha_s = 110$ *dB (solid line in both plots),* $N = 4$*,* $M = 8$ *and* $\alpha_s = 220$ *dB (dashed line in* (b)*), and* $N = 8$*,* $M = 8$ *and* $\alpha_s = 110$ *dB (dashed line in* (c)*).*

difference measure between the look-up table amplitude response and $D(\omega)$ in a set of designer-defined subbands. The difference measure is usually frequency-dependent, i.e., it can be expressed as a weighted error function given by

$$E(\omega) = W(\omega)\left(G(\omega) - D(\omega)\right), \tag{23}$$

where $G(\omega)$ is the amplitude response of the look-up table and $W(\omega)$ is a non-negative frequency-dependent weighting function. By the choice of the error measure, different look-up-table designs are obtained. By minimizing the maximum absolute value of $E(\omega)$ in the subbands of interest R,

$$\epsilon_{\max} = \max_{\omega \in R} |E(\omega)|, \tag{24}$$

a measure called the Chebyshev or minimax criterion or L^∞ norm, the look-up table can designed using the Parks-McClellan algorithm [36]. An alternative is to minimize the integral of the squared error function

$$\epsilon_{\text{ls}} = \int_{\omega \in R} E^2(\omega)\mathrm{d}\omega, \tag{25}$$

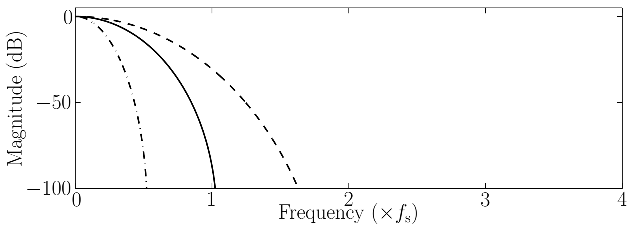

Figure 4: *Effect of the stopband attenuation and the number of sample to be corrected on the alias reduction performance with the Dolph-Chebyshev window for table parameters* $N = 4$*,* $M = 8$*, and* $\alpha_s = 110$ *dB (solid line),* $N = 4$*,* $M = 8$ *and* $\alpha_s = 220$ *dB (dashed line), and* $N = 8$*,* $M = 8$ *and* $\alpha_s = 110$ *dB (dash-dotted line).*

a measure called the least-squares criterion, or minimizing the L^2 norm.

In the previous sections it was observed that the aliasing is traded off for amplitude attenuation at high frequencies. Since the relation becomes tighter as the look-up-table size and especially N is reduced, the amplitude attenuation is unavoidable when a small look-up table is used. However, the attenuation amount should be limited to a certain level because the compensation for an excessive amplitude drop may complicate the post-processing or boost the aliased components back to their original level. Therefore, an appropriate attenuation level can be specified at a chosen high frequency as a constraint for optimization. This can be regarded as the constraint on the passband given by

$$1 - \alpha_{\text{pb}} \leq G(\omega) \leq 1 \tag{26}$$

for $\omega \in [0, \omega_{\text{pb}}]$, where α_{pb} denotes the maximum allowed amplitude attenuation at ω_{pb}. Once this condition is satisfied for the passband, the filter design can focus on minimizing the stopband level in order to suppress the aliased components. Since the desired amplitude response at the stopband is zero, the minimization objective is given as

$$\min ||W(\omega)G(\omega)||_p \tag{27}$$

for $\omega \in [\omega_{\text{sb}}, \pi]$, where p gives the minimax criterion when $p = \infty$, or the least squares criterion when $p = 2$.

Furthermore, the amplitude response $G(\omega)$ can be represented in terms of the impulse response $g(k)$ which is the set of variables to be optimized. The frequency response corresponding to $g(k)$ on the frequency grid ω_n is given by

$$H_g(\omega_n) = \sum_{k=-NM/2}^{NM/2} g(k)e^{-j\omega_n k} \tag{28}$$

$$= g(0) + 2\sum_{k=1}^{NM/2} g(k)\cos(\omega_n k). \tag{29}$$

Using this, Eqs. (26) and (27) can be expressed in matrix form as

$$1-\alpha_{\mathrm{pb}} \preceq \begin{bmatrix} 1 & 2\cos(\omega_0) & \dots & 2\cos(\omega_0 L) \\ \vdots & & & \vdots \\ 1 & 2\cos(\omega_{\mathrm{pb}}) & \dots & 2\cos(\omega_{\mathrm{pb}} L) \end{bmatrix} g \preceq 1, \quad (30)$$

$$\min \left\| \mathrm{diag}(W) \begin{bmatrix} 1 & 2\cos(\omega_{\mathrm{sb}}) & \dots & 2\cos(\omega_{\mathrm{sb}} L) \\ \vdots & & & \vdots \\ 1 & 2\cos(\pi) & \dots & 2\cos(\pi L) \end{bmatrix} g \right\|_p, \quad (31)$$

respectively, where $\mathrm{diag}(W)$ denotes a diagonal matrix with the weighting function values $W(\omega_n)$ at the chosen stopband frequencies as its diagonal elements. Note that $G(\omega)$ is equal to $H_g(\omega)$ over the passband due to the zero phase. Finally, the optimization problem is set up as

$$\min ||\mathrm{diag}(W) F_{\mathrm{sb}} g||_p \quad (32)$$

$$\text{subject to } F_{\mathrm{pb}} g \preceq 1, \quad (33)$$

$$F_{\mathrm{pb}} g \succeq 1-\alpha_{\mathrm{pb}}, \text{ and} \quad (34)$$

$$F_{\mathrm{dc}} g = 1, \quad (35)$$

where F_{pb} and F_{sb} are the Fourier transform matrices in Eqs. (30) and (31) and F_{dc} is the first row of F_{pb} for frequency 0. The last constraint is added to make the sum of the impulse response be equal to one. This optimization problem can be solved by linear programming when $p = \infty$ or by quadratic programming when $p = 2$, and there are efficient algorithms for both cases [32].

This optimization approach is exemplified in Fig. 5(a) with an equal weighting, i.e., $W(\omega) = 1$ for all ω, using a passband criterion $\alpha_{\mathrm{pb}} = 6$ dB at 15 kHz, a stopband from $f_{\mathrm{s}}-$ 5 kHz to $Mf_{\mathrm{s}}/2$, and table-size parameters $N = 4$ and $M = 8$. It can be seen that the L^2 norm provides a better rejection over the stopband than the L^∞ norm, at the expense of the slightly wider passband. In practice, these diffences are small, as illustrated in Figs. 5(b) and 5(c), where the spectra of a sawtooth waveform synthesized with the BLIT algorithm using these look-up tables are shown. However, the aliasing level below the fundamental frequency can be observed to be lower in the L^2 norm approach, although the first-order generation of aliased components is slightly higher.

As one can observe, the stopband attenuation of the optimization approach is slightly worse than with the parametric-window approach. This results in greater aliasing in general. However, since human hearing is more sensitive at low and middle frequencies, aliasing at high frequencies can be more acceptable. Therefore, the filter optimization can incorporate a weighting function $W(\omega)$ that gives more weight to the frequencies that will fold back to the low and middle frequencies when the look-up table is downsampled at the synthesis stage. An idea similar to using such a weighting in optimal filter design was demonstrated in the design of an audio interpolator for the purpose of pitch-shifting in sampling synthesis [37]. The main technique was to place stopband notches on multiples of the sampling frequency.

The frequency-dependent weighting approach is illustrated in Fig. 6. A discrete-valued weighting function is chosen to be 1 over all frequency ranges that will alias to the interval [0,10] kHz, including that interval itself, and zero otherwise. The weighting function was obtained by multiplying the discrete-valued function with $\lceil f/f_{\mathrm{s}} \rceil^{\rho}$ to adjust the spectral tilt over the stopband by imposing more weight for higher frequencies. The parameter ρ controls the tilt, and it was set to 3.2 in this example. The resulting stopband is relaxed over the frequencies that fold back between 10 kHz and the Nyquist frequency, whereas the stopband corresponding to low frequencies is more suppressed than in the equally-weighted example. Furthermore, the additional weighting factor effectively tilts the overall stopband downward.

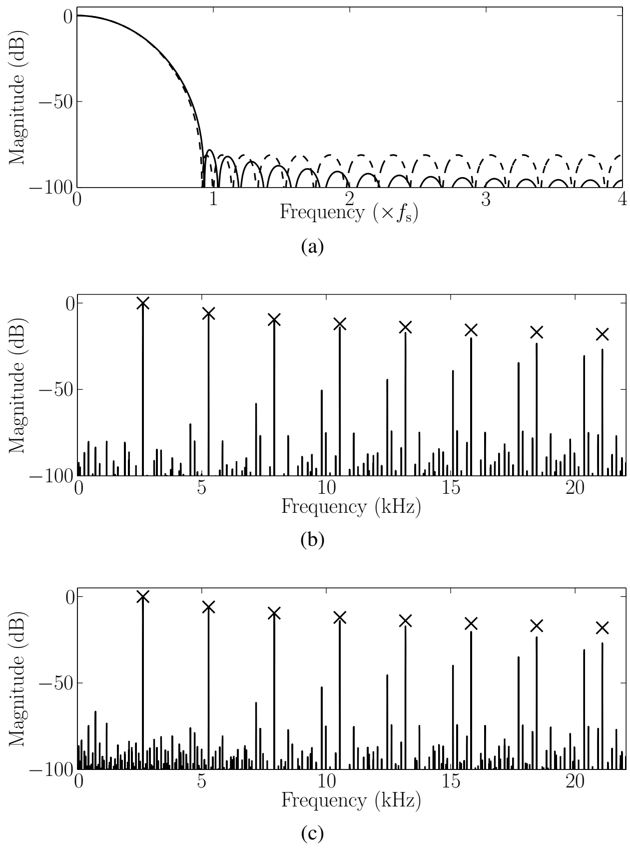

Figure 5: (a) *Magnitude responses of the equal-weighted look-up table designs that minimize the L^2 norm (solid line) and L^∞ norm (dashed line) of the error function using table-size parameters $N = 4$, $M = 8$, a passband criterion $\alpha_{\mathrm{pb}} = 6$ dB at 15 kHz, and a stopband from $f_{\mathrm{s}}-$ 5 kHz to $Mf_{\mathrm{s}}/2$. The spectra of the sawtooth waveforms synthesized with the BLIT algorithm using L^2- and L^∞-optimized look-up tables are shown in* (b) *and* (c)*, respectively.*

The improved aliasing reduction for the unequally-weighted example is illustrated in Figs. 6(b) and 6(c), with the spectra of the sawtooth waveforms having been synthesized with the BLIT algorithm using look-up tables designed as described. The L^∞-optimized look-up table is significantly better than the equally-weighted counterpart, having significantly less aliasing below 10 kHz. Moreover, no qualitative difference can be seen between the magnitude spectra. Furthermore, the degree of aliasing reduction is approximately similar to that of the parametric-window approaches—see Figs. 3 and 4.

5. CONCLUSIONS AND FURTHER WORK

In quasi-bandlimited oscillator algorithms, the aliasing corrupting the trivially sampled classical geometric waveforms, such as the sawtooth wave, is reduced by applying a lowpass filter prior to

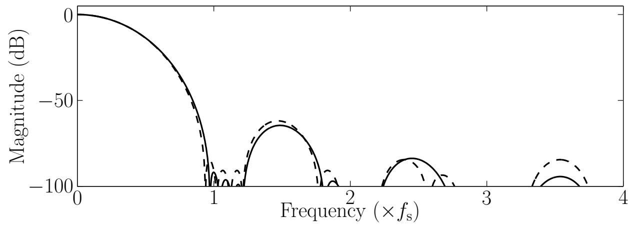

(a)

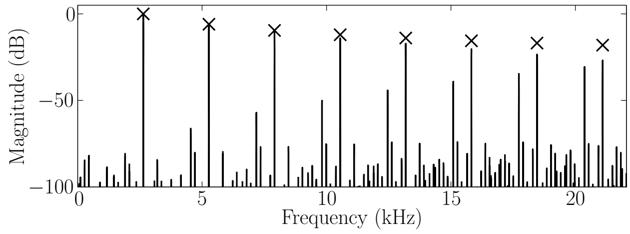

(b)

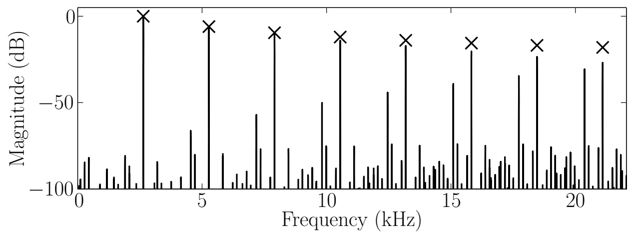

(c)

Figure 6: (a) *Magnitude responses of the weighted look-up-table designs that minimize the L^2 norm (solid line) and L^∞ norm (dashed line) of the error function using parameters $N = 4$, $M = 8$. The spectra of the sawtooth waveforms using L^2- and L^∞-optimized look-up tables are shown in* (b) *and* (c)*, respectively.*

sampling of the ideal waveform. Effectively the same result can be obtained in the digital domain by applying a correction function that replaces the non-bandlimited discontinuities in the waveform or its derivative by a bandlimited counterpart. In the ideal case the correction function can be expressed in a closed form, but the ideal correction function is infinitely long and computationally inefficient to evaluate for every sample. Therefore, the ideal correction function is typically windowed and tabulated in a look-up table.

In order to have a good alias-reduction performance, the windowed ideal correction function still needs to be quite large. It was shown in this paper that when a short look-up table is used, the windowed sinc function is *not* perceptually optimal. In particular, less audible aliasing is obtained, at the price of less objectionable spectral rolloff, by sampling the window function alone in place of a windowed sinc.

A perceptually optimal look-up table design was formulated as the solution to a highly nonlinear and implicit optimization problem. However, since the existence of a feasible solution to that optimization is not guaranteed, practical parametric look-up table designs were discussed and evaluated in this paper. Of the suggested designs, the parametric window functions (Dolph-Chebyshev and Kaiser windows) and the least-squares optimized multi-band FIR-filter designs provided the best alias reduction performance. They are therefore recommendable approaches to look-up table design.

The choice of the frequency-dependent weighting function in the constraint-based multi-band FIR-filter design affects greatly the alias-reduction performance of the resulting look-up table. In this paper, only simplified weighting functions were tried, and a weighting function that matches more closely the aliasing audibility threshold could provide better performance. Therefore, investigations into perceptually optimal weighting function should be carried out.

Furthermore, as the optimality of the proposed look-up table designs was not tested in this paper, perceptually optimal parameters should be found for the proposed designs at various fundamental frequencies. These parameters could then be used to design look-up tables yielding perceptually optimal aliasing reduction. In addition, with perceptually optimal parameters, the proposed look-up table designs could be analyzed in more detail, and they could be compared to other existing antialiasing oscillator algorithms.

The look-up tables and sound examples for the cases studied in this paper can be found at `http://www.acoustics.hut.fi/go/dafx10-optosctables/`.

6. REFERENCES

[1] H. F. Olson and H. Belar, "Electronic music synthesis," *Journal of the Acoustical Society of America*, vol. 27, no. 3, pp. 595–612, May 1955.

[2] R. A. Moog, "Voltage-controlled electronic music modules," *Journal of the Audio Engineering Society*, vol. 13, no. 3, pp. 200–206, July 1965.

[3] A. Chaudhary, "Bandlimited simulation of analog synthesizer modules by additive synthesis," in *Proceedings of the Audio Engineering Society 105th Convention*, San Francisco, CA, September 1998, preprint number 4779.

[4] G. Winham and K. Steiglitz, "Input generators for digital sound synthesis," *Journal of the Acoustical Society of America*, vol. 47, part 2, pp. 665–666, February 1970.

[5] J. A. Moorer, "The synthesis of complex audio spectra by means of discrete summation formulae," *Journal of the Audio Engineering Society*, vol. 24, no. 9, pp. 717–727, November 1976.

[6] G. Deslauriers and C. Leider, "A bandlimited oscillator by frequency-domain synthesis for virtual analog applications," in *Proceedings of the Audio Engineering Society 127th Convention*, New York, NY, October 2009, preprint number 7923.

[7] P. Burk, "Band limited oscillators using wave table synthesis," in *Audio Anecdotes II — Tools, Tips, and Techniques for Digital Audio*, K. Greenebaum and R. Barzel, Eds., pp. 37–53. A. K. Peters Ltd., Wellesley, MA, 2004.

[8] J. Nam, V. Välimäki, J. S. Abel, and J. O. Smith, "Alias-free oscillators using feedback delay loops," in *Proceedings of the 12th International Conference on Digital Audio Effects (DAFx-09)*, Como, Italy, September 2009, pp. 347–352.

[9] E. Brandt, "Hard sync without aliasing," in *Proceedings of the International Computer Music Conference*, Havana, Cuba, September 2001, pp. 365–368.

[10] T. S. Stilson and J. O. Smith, "Alias-free digital synthesis of classic analog waveforms," in *Proceedings of the International Computer Music Conference*, Hong Kong, China, August 1996, pp. 332–335.

[11] T. S. Stilson, *Efficiently-Variable Non-Oversampling Algorithms in Virtual-Analog Music Synthesis – A Root-Locus Perspective*, Ph.D. thesis, Department of Electrical Engineering, Stanford University, Stanford, CA, June 2006, available online `http://ccrma.stanford.edu/~stilti/papers/` [Checked June 25, 2010].

[12] J. Timoney, V. Lazzarini, and T. Lysaght, "A modified FM synthesis approach to bandlimited signal generation," in *Proceedings of the 11th International Conference on Digital Audio Effects (DAFx-08)*, Espoo, Finland, September 2008, pp. 27–33.

[13] J. Nam, V. Välimäki, J. S. Abel, and J. O. Smith, "Efficient antialiasing oscillator algorithms using low-order fractional delay filters," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 4, pp. 773–785, May 2010.

[14] J. Pekonen, V. Välimäki, J. Nam, J. S. Abel, and J. O. Smith, "Variable fractional delay filters in bandlimited oscillator algorithms for music synthesis," in *Proceedings of the 2010 International Conference on Green Circuits and Systems (ICGCS2010)*, Shanghai, China, June 21–23 2010, pp. 148–153.

[15] H. Chamberlin, *Musical Applications of Microprocessors*, pp. 418–424, Hayden Book Company, Hasbrouck Heights, NJ, 2nd edition, 1985.

[16] M. Puckette, *The Theory and Technique of Electronic Music*, pp. 301–322, World Scientific Publishing Co., Hackensack, NJ, 2007.

[17] V. Välimäki, "Discrete-time synthesis of the sawtooth waveform with reduced aliasing," *IEEE Signal Processing Letters*, vol. 12, no. 3, pp. 214–217, March 2005.

[18] A. Huovilainen and V. Välimäki, "New approaches to digital subtractive synthesis," in *Proceedings of the International Computer Music Conference*, Barcelona, Spain, September 2005, pp. 399–402.

[19] V. Välimäki and A. Huovilainen, "Oscillator and filter algorithms for virtual analog synthesis," *Computer Music Journal*, vol. 30, no. 2, pp. 19–31, Summer 2006.

[20] V. Välimäki, J. Nam, J. O. Smith, and J. S. Abel, "Alias-suppressed oscillators based on differentiated polynomial waveform," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 4, pp. 786–798, May 2010.

[21] J. Lane, D. Hoory, E. Martinez, and P. Wang, "Modeling analog synthesis with DSPs," *Computer Music Journal*, vol. 21, no. 4, pp. 23–41, Winter 1997.

[22] J. Pekonen, "Computationally efficient music synthesis – methods and sound design," Master of Science (Technology) thesis, TKK Helsinki University of Technology, Espoo, Finland, June 2007, available online `http://www.acoustics.hut.fi/publications/files/theses/pekonen_mst/` [Checked June 25, 2010].

[23] J. Kleimola, "Audio synthesis by bitwise logical modulation," in *Proceedings of the 11th International Conference on Digital Audio Effects (DAFx-08)*, Espoo, Finland, September 2008, pp. 67–70.

[24] J. Timoney, V. Lazzarini, J. Pekonen, and V. Välimäki, "Spectrally rich phase distortion sound synthesis using an allpass filter," in *Proceedings of the 2009 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'09)*, Taipei, Taiwan, April 2009, pp. 293–296.

[25] J. Timoney, V. Lazzarini, B. Carty, and J. Pekonen, "Phase and amplitude distortion methods for digital synthesis of classic analogue waveforms," in *Proceedings of the Audio Engineering Society 126th Convention*, Munich, Germany, May 2009, preprint number 7792.

[26] J. Kleimola, J. Pekonen, H. Penttinen, V. Välimäki, and J. S. Abel, "Sound synthesis using an allpass filter chain with audio-rate coefficient modulation," in *Proceedings of the 12th International Conference on Digital Audio Effects (DAFx-09)*, Como, Italy, September 2009, pp. 305–312.

[27] V. Lazzarini, J. Timoney, J. Kleimola, and V. Välimäki, "Five variations on a feedback theme," in *Proceedings of the 12th International Conference on Digital Audio Effects (DAFx-09)*, Como, Italy, September 2009, pp. 139–145.

[28] V. Lazzarini and J. Timoney, "New perspectives on distortion synthesis for virtual analog oscillators," *Computer Music Journal*, vol. 34, no. 1, pp. 28–40, Spring 2010.

[29] J. Pekonen and V. Välimäki, "Filter-based alias reduction for digital classical waveform synthesis," in *Proceedings of the 2008 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'08)*, Las Vegas, NV, April 2008, pp. 133–136.

[30] V. Välimäki and A. Huovilainen, "Antialiasing oscillators in subtractive synthesis," *IEEE Signal Processing Magazine*, vol. 24, no. 2, pp. 116–125, March 2007.

[31] A. B. Leary and C. T. Bright, "Bandlimited digital synthesis of analog waveforms," U.S. Patent 7,589,272, September 2009.

[32] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty, *Nonlinear Programming: Theory and Algorithms*, chapter 9, pp. 469–536, John Wiley & Sons, Inc., Hoboken, NJ, 3rd edition, 2006.

[33] J. F. Kaiser, "Nonrecursive digital filter design using the I_0-sinh window function," in *Proceedings of the 1974 IEEE International Symposium on Circuits and Systems*, San Francisco, CA, April 1974, pp. 20–23.

[34] H. D. Helms, "Nonrecursive digital filters: Design methods for achieving specifications on frequency response," *IEEE Transactions on Audio and Electroacoustics*, vol. 16, no. 3, pp. 336–342, September 1968.

[35] E. Zwicker and H. Fastl, *Psychoacoustics*, pp. 15–19, Springer-Verlag, Berlin, Germany, 1990.

[36] T. W. Parks and J. H. McClellan, "Chebyshev approximation for nonrecursive digital filters with linear phase," *IEEE Transactions on Circuit Theory*, vol. 19, no. 2, pp. 189–194, March 1972.

[37] D. Rossum, "Constraint based audio interpolators," in *Proceedings of the 1993 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA'93)*, New Paltz, NY, October 1993, pp. 161–164.