FILTER-BASED ALIAS REDUCTION FOR DIGITAL CLASSICAL WAVEFORM SYNTHESIS

Jussi Pekonen and Vesa Välimäki

TKK Helsinki University of Technology
Department of Signal Processing and Acoustics
[jussi.pekonen|vesa.valimaki]@tkk.fi

ABSTRACT

The classical waveforms used in the subtractive sound synthesis have rich spectral content, which causes their sampled digital implementations to suffer from aliasing distortion. Several antialiasing waveform synthesis algorithms have been suggested, and they either remove the aliasing completely or reduce it greatly. A new approach to alias reduction is proposed where the remaining aliased components are suppressed by applying digital highpass and/or comb filtering to the output of an antialiasing algorithm. Applicable filter designs for this novel postprocessing approach are discussed and evaluated with respect to the alias reduction performance using noise-to-mask ratio (NMR). The NMR can be reduced by 10 dB at high fundamental frequencies with a computationally efficient highpass filter. The NMR can be further reduced by using a combination of an IIR comb filter and a DC blocking filter, which provides the best alias reduction performance at high fundamental frequencies.

Index Terms— Acoustic signal processing, antialiasing, audio oscillators, music, signal synthesis.

1. INTRODUCTION

The subtractive sound synthesis generates sound by filtering a spectrally rich waveform, such as the sawtooth and the rectangular wave [1]. These classical continuous-time waveforms contain infinitely many harmonically related frequencies, which can be seen from their Fourier series representations. Therefore, the trivial sampling of the continuous-time waveforms suffers from aliasing, since the frequencies above the Nyquist limit are folded on to the audio band. Especially at high fundamental frequencies, this aliasing distortion becomes clearly audible and disturbing.

Human hearing is a highly nonlinear process, and the frequency resolution of hearing makes no exception [2]. The frequencies of the audio band are perceived approximately logarithmically so that the middle point of the perceived audio band is about 1 kHz. Due to this behavior, the aliased components below the fundamental frequency of a waveform are easily perceived at high fundamental frequencies. At lower fundamental frequencies the aliasing between the waveform harmonics becomes more disturbing. Although the harmonics of a waveform cause some of the aliased components to become inaudible due to the phenomenon of frequency masking, not all of them are masked.

In this paper, the task of alias reduction in digital classical waveform synthesis using postprocessing digital filtering is investigated. Existing antialiasing waveform synthesis algorithms are briefly reviewed in Section 2. In Sections 3 and 4, two approaches for alias reduction in trivially sampled continuous-time waveform and antialiasing algorithms are presented. The evaluation of the presented approaches with respect to the alias reduction performance is presented in Section 5. Section 6 concludes the paper.

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

2. ANTIALIASING OSCILLATOR ALGORITHMS

There are several antialiasing classical waveform oscillator algorithms, which have been previously classified into three categories as follows [3]:

1. bandlimited algorithms, where only frequencies below the Nyquist limit are generated,
2. quasi-bandlimited algorithms, where a continuous-time waveform is first lowpass filtered and then sampled, and
3. alias-suppressing algorithms, where the spectral tilt of the waveform is modified.

The first class contains algorithms which are ideally bandlimited, i.e., they contain no aliasing. This class includes, e.g., the additive synthesis technique [4]. The algorithms of the second class allow some aliasing mainly at high frequencies, and they are implemented by performing antialiasing filtering to the continuous-time waveform before sampling in order to attenuate the frequencies above the Nyquist limit. With simple geometric waveforms, such as the classical waveforms, this can be performed in closed form, e.g., by integrating a bandlimited impulse train [5, 6]. The third class contains algorithms which allow aliasing in the whole audio band but sufficiently suppressed. These methods usually apply an exponential decay function to the harmonic amplitudes before sampling and restore the original spectral tilt with a digital filter. This class includes, e.g., the differentiated parabolic waveform (DPW) technique [7, 8].

The postprocessing approaches proposed in this paper do not fulfill the definition of any of the abovementioned categories, and thus they form a fourth category. This new class, filter-based postprocessing algorithms, applies digital filtering to a waveform which is obtained by either trivial sampling or as the output of an antialiasing algorithm of categories two or three.

3. HIGHPASS FILTERING APPROACH

Since the aliasing below the fundamental frequency is the dominant reason for the alias disturbance at high fundamental frequencies, some of the aliasing distortion can be reduced by applying highpass filtering with which the components below the fundamental frequency are attenuated. Next, practical filter designs for this highpass filtering approach are discussed.

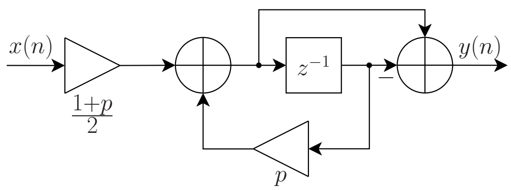

Fig. 1. DC blocking filter.

Since the fundamental frequencies in music are usually only small fractions of the sampling rate, the transition band of the filter becomes narrow. This leads to large filter orders in traditional filter designs, which is often not desirable due to the reduced computational efficiency. The problem of a narrow transition band can be relaxed by interpolating a prototype filter designed at a lower sampling rate [9] or obtained by a spectral transformation [10], yielding improved computational efficiency. However, especially in real-time applications, the fundamental frequency of the waveform is time-varying, and, in order to obtain the maximal alias reduction performance, the filter coefficients should be updated fast, preferably on a per-sample basis [11]. The coefficient value update in traditional filter designs requires the computation of each coefficient separately, which decreases the computational efficiency of the algorithm. Therefore, the traditional filter designs are impractical in this approach, and alternative designs must be investigated.

A computationally efficient and easily adjustable highpass filter is obtained by setting the filter to have a transmission zero at the zero frequency and by using a pole to set the amplification at other frequencies. This first order IIR filter, the DC blocker, is given by

$$H_{\mathrm{DC}}(z) = \frac{1+p}{2}\frac{1-z^{-1}}{1-pz^{-1}}, \tag{1}$$

where $|p| < 1$, and its block diagram is shown in Figure 1. The factor $(1+p)/2$ is for normalization, as it sets the gain at the Nyquist limit to unity. The filter structure requires only two multiplications and two additions per sample, and the multipliers are easily updated by changing the pole p and computing the normalization factor.

The pole can be calculated from the desired -3 dB frequency $f_{\mathrm{c}} \in (0, f_{\mathrm{s}}/2)$ by solving p from equation $|H_{\mathrm{DC}}(e^{j2\pi f_{\mathrm{c}}/f_{\mathrm{s}}})|^2 = 1/2$, yielding

$$p = \frac{1-\sin\left(\frac{2\pi f_{\mathrm{c}}}{f_{\mathrm{s}}}\right)}{\cos\left(\frac{2\pi f_{\mathrm{c}}}{f_{\mathrm{s}}}\right)} = \tan\left(\frac{\pi}{4} - \frac{\pi f_{\mathrm{c}}}{f_{\mathrm{s}}}\right), \tag{2}$$

where f_{s} is the sampling frequency. The latter form is obtained by using the properties of trigonometric functions. However, computation of the exact pole position requires evaluation of a trigonometric function, which makes the accurate pole calculation computationally inefficient. In order to compute the pole position efficiently, (2) must be approximated with a low order polynomial. A computationally efficient first order approximation is given by

$$\hat{p} = 1 - \frac{2}{\pi}\frac{2\pi f_{\mathrm{c}}}{f_{\mathrm{s}}} = 1 - \frac{4f_{\mathrm{c}}}{f_{\mathrm{s}}}, \tag{3}$$

which produces slightly lower and higher -3 dB frequencies than (2) below and above $f_{\mathrm{s}}/4$, respectively. Note that the computation of the pole position using (3) requires no division, since $1/f_{\mathrm{s}}$ is a constant.

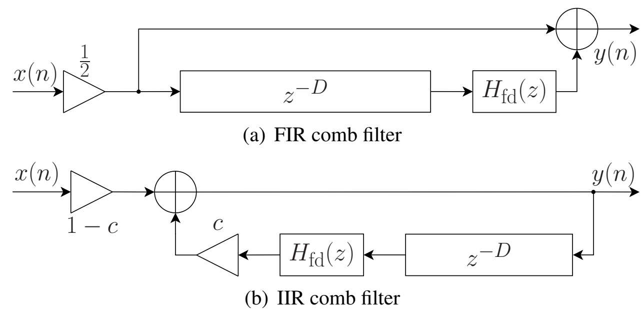

Fig. 2. FIR and IIR comb filters.

The DC blocking filter can be computationally more efficient, if the filter normalization is omitted. Since the fundamental frequencies found in music are rarely above 4200 Hz, the filter provides less than 2 dB amplification at the Nyquist limit below the 4200 Hz limit without normalization when the sampling frequency is above 40 kHz. Therefore, the normalization multiplication can be omitted.

4. COMB FILTERING APPROACH

Since at low fundamental frequencies the aliasing between the waveform harmonics becomes more disturbing, the level of the aliased components in between the harmonics can be decreased by using a comb filter, which passes harmonically related frequencies and attenuates the rest. Next, filter designs for the comb filtering approach are discussed.

A comb filter can be constructed in two alternative ways [12]. The first variation, the FIR comb filter, passes harmonically related frequencies and attenuates the frequencies in between them. The second variation, the IIR comb filter, applies large amplification to harmonically related frequencies and smaller amplification to the frequencies in between them. The transfer functions of gain-normalized FIR and IIR comb filters are given by

$$H_{\mathrm{FIR}}(z) = \frac{1}{2}\left(1 + z^{-L}\right), \tag{4}$$

$$H_{\mathrm{IIR}}(z) = \frac{1-c}{1-cz^{-L}}, \tag{5}$$

respectively. The feedback coefficient $c \in (0, 1)$ in the IIR comb filter sets the attenuation between the harmonics. The delay-line length L defines the frequencies to be passed by the comb filter, and for a waveform with fundamental frequency f_0 it is given by $L = f_{\mathrm{s}}/f_0$. Thus, when the fundamental frequency f_0 is varied, a division is required to update the delay-line length L.

Since the fundamental frequency of a waveform can be arbitrary, the required delay-line length L is rarely an integer. Therefore, in order to avoid the attenuation of the waveform harmonics, the fractional part of the desired delay must be implemented using a fractional delay filter. The use of the fractional delay filter $H_{\mathrm{fd}}(z)$ is indicated in Figure 2, where the block diagrams of the FIR and IIR comb filters are illustrated. In Figure 2, D denotes the integer part of L.

In practice, the fractional delay is implemented using either a first order Lagrange (linear, FIR) interpolator or a first order Thiran allpass filter [13]. However, both of these implementations are dispersive, i.e., the produced delay is not constant for all frequencies.

At low frequencies the delay approximates quite well the target, but at high frequencies it deviates from the desired depending on the implementation type and the targeted delay. However, since the frequency resolution of the human hearing is approximately logarithmic, the deviation of the higher harmonics is practically inaudible.

In addition, since the linear interpolator is in practice a lowpass filter, it attenuates the higher harmonics of the waveform. The attenuation is more severe in the IIR comb filter, tens of decibels, while in the FIR comb filter the attenuation is only a few decibels. However, since in the subtractive sound synthesis the source waveform is usually filtered with a lowpass filter, the lowpass characteristics of the linear interpolator can be taken into account when designing the synthesizer filter.

When the FIR and IIR comb filters are defined as in (4) and (5), respectively, they also pass frequencies near the zero frequency. These aliased components can be attenuated by applying the DC blocking filter described in Section 3. However, if the desired waveform has only odd harmonics, as in the case of the triangular pulse wave or the rectangular pulse wave with a duty cycle of 50%, the DC blocking filter is not needed if the addition in the comb filter is replaced with subtraction and the delay line length is halved, i.e., $L = f_{\mathrm{s}}/(2f_0)$.

5. EVALUATION OF THE FILTERING APPROACHES

The alias reduction performances of the proposed filtering methods are evaluated using the noise-to-mask ratio (NMR), a measure proposed for evaluation of perceptual audio coding methods [14, 15]. The NMR figures are computed in a similar manner as in [3], and the algorithm is briefly reviewed next.

The NMR algorithm requires two signals, a bandlimited reference signal and the difference between the alias corrupted signal and the reference signal. The amplitude and phase errors in the evaluation are avoided by composing the reference signal from estimated amplitudes and phases of the harmonics from the corrupted signal. Next, a 1024-point magnitude spectrum of the reference and error signals is computed with the FFT using the Hann window, and the spectra are divided into segments which approximate the critical bands of hearing. The average energy of each band is scaled by dividing by its width in bins and the average energy is converted to the decibel scale. The frequency masking phenomenon is simulated by copying a fraction of the signal energy from each critical band to all neighboring bands using an interband spreading function. The hearing threshold is applied as an additive term, the final energy per band is computed, and energies of all bands are added up. Finally, a single NMR figure is obtained as the ratio of the error to the mask threshold. In the evaluation, the sampling rate of 44.1 kHz was used. This affects the choice of critical bandwidths and spreading functions. Smaller NMR values are considered better, and the NMR values below -10 dB are considered to be free from audible artifacts [15].

In Figure 3, the NMRs of the trivially sampled sawtooth wave and its highpass and comb filtered version are presented for 88 fundamental frequencies spanning the piano range from 27.5 Hz to 4186 Hz. The NMR of the DPW algorithm [7] is also plotted for comparison. The fractional delays in FIR and IIR comb filters are implemented using the linear and the first order allpass interpolators, respectively. The feedback coefficient of the IIR comb filter is $c = 0.995$, which yields approximately 52 dB attenuation of the frequencies between the harmonics. The outputs of the comb filters are also filtered with the DC blocking filter. The pole of the DC blocking filter is calculated with (3) using $f_{\mathrm{c}} = f_0$ in all cases.

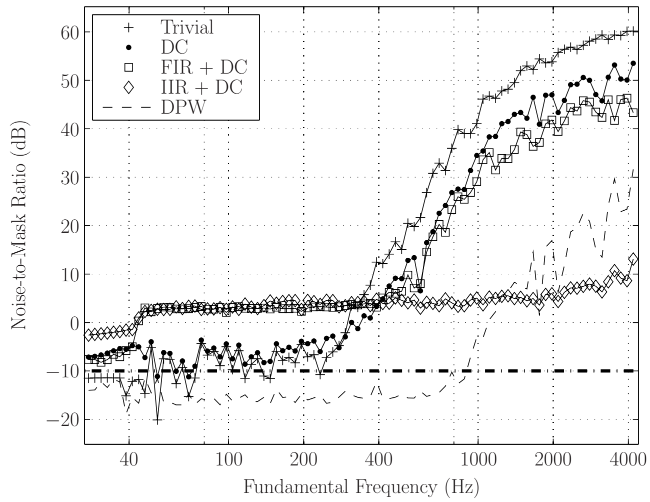

Fig. 3. The NMR figures of a trivially sampled sawtooth wave and its highpass and comb-filtered versions. The NMR of the DPW algorithm (dashed line) is plotted for comparison.

It is seen in Figure 3 that above 350 Hz all proposed filtering methods perform better than the trivial sampling. The DC blocking filter provides approximately 10 dB and the FIR comb filter in combination with the DC blocking filter approximately 15 dB improvements. The IIR comb filter with the DC blocking filter provides an almost flat 6 dB NMR figure for all tested fundamental frequencies. Only the IIR approach provides improvement over the DPW algorithm above 1100 Hz, while the DC blocking and FIR approaches provide NMR figures between the trivial sampling and the DPW.

Below 350 Hz the DC blocking filter provides approximately the same NMR as the trivial sampling, and the FIR comb filter with the DC blocking filter provides the same 6 dB NMR figure as the IIR comb filter. The larger NMR figures of the comb filtering approaches are due to the dispersion of the fractional delay, which causes the higher harmonics to be attenuated while some aliased components near the harmonic frequencies are passed without any attenuation. Therefore, the waveform contains effectively more aliasing at high frequencies, which increases the NMR figure of the comb filtering approaches at low fundamental frequencies.

In Figure 4, the NMRs of the DPW algorithm and its filtered versions are presented using the same filter parameters as for the trivially sampled sawtooth wave. Now, the NMR of the DPW2X, a multirate version of the DPW algorithm [7], is plotted for comparison. The proposed filtering approaches provide improvements to the DPW algorithm above 1100 Hz, below which all approaches produce larger NMR figures. With the comb filtering approaches they are again due to the fractional delay dispersion, but the worse performance of the DC blocking filter is due to the properties of the DPW algorithm. The DPW algorithm produces less aliasing at low frequencies than the trivial sampling, and at low fundamental frequencies the DC blocking filter attenuates the lowest harmonics, which cause the signal to contain more aliasing at high frequencies.

It is seen in Figure 4 that above 1100 Hz the DC blocking filter provides approximately 7 dB improvement, and the FIR comb filter with the DC blocking filter provides approximately 12 dB improvement. Yet again, the IIR comb filter in combination with the DC blocking filter provides the best performance, an approximately constant 2 dB NMR figure for all fundamental frequencies. Above

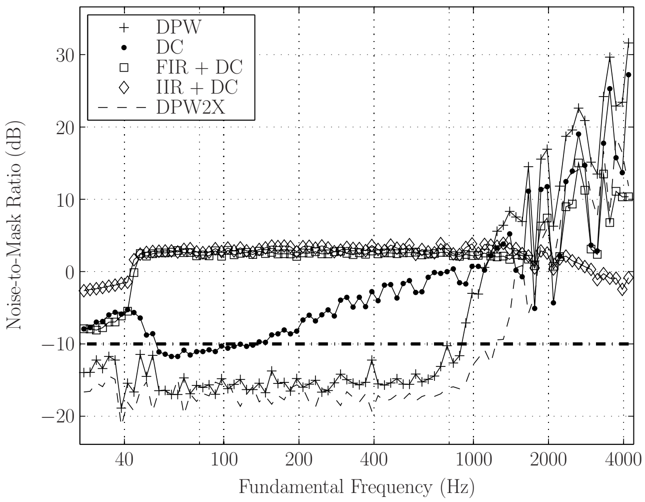

Fig. 4. The NMR figures of the DPW algorithm and its highpass and comb-filtered versions. The NMR of the DPW2X algorithm (dashed line) is plotted for comparison.

Table 1. The additional multiplications (MPY) and additions (ADD) required in the filtering approaches and the DPW algorithms with respect to the trivial sampling approach (one modulo counter).

Method	MPY	ADD	*Note*
DC block	2 (1)	2	No division
FIR comb	1	1	Requires a division
IIR comb	2	1	Requires a division
DPW	2	1	Requires a division
DPW2X	3	3	Additional modulo counter

1750 Hz, the DC blocking filter and the FIR comb filter provides approximately the same performance as the DPW2X algorithm, while the IIR comb filter provides better performance.

In Table 1, the additional multiplication and addition operations compared to the trivial sampling approach required in each presented approach are listed. Although the proposed filtering methods require a few additional operations in order to be comparable with the DPW and DPW2X algorithms, they do provide better performance at high fundamental frequencies. Especially, if the division is not desirable, the proposed DC blocking filter provides an intermediate solution in terms of alias reduction performance and computational efficiency.

If a better alias reduction performance is desired, the algorithm requires the division. By using the DPW algorithm in combination with the IIR comb filter and the DC blocking filter, the aliasing can be reduced greatly at high fundamental frequencies. At low frequencies the postprocessing filters, or the comb filter alone, could be switched off, providing approximately the same alias reduction performance as the original algorithm.

6. CONCLUSIONS

Three different postprocessing digital filtering approaches for the task of alias reduction of digital classical waveform synthesis algorithms were proposed. The NMR was used in the evaluation of the proposed algorithms. All approaches provide improvements over the existing antialiasing algorithms in alias reduction at high fundamental frequencies with only a slight increase in the computational complexity. The NMR can be decreased by 10 dB by applying the proposed computationally efficient highpass filter, which requires no division. At high fundamental frequencies, the proposed combination of an IIR comb filter and a DC blocking filter provides the best alias reduction performance among the considered techniques.

7. REFERENCES

[1] C. Roads, *The Computer Music Tutorial*, The MIT Press, Cambridge, Massachusetts, USA, 1995.

[2] J. O. Pickles, *An Introduction to the Physiology of Hearing*, Academic Press, London, UK, 2nd edition, 1988.

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

[4] A. Chaudhary, “Band-limited simulation of analog synthesizer modules by additive synthesis,” in *Proceedings of the AES 105th Convention*, San Francisco, California, USA, August 1998.

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

[6] T. S. Stilson, *Efficiently-Variable Non-Oversampled Algorithms in Virtual-Analog Music Synthesis – A Root-Locus Perspective*, Ph.D. thesis, Stanford University, Stanford, California, USA, 2006.

[7] 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.

[8] 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.

[9] Y. Neuvo, D. Cheng-Yu, and S. K. Mitra, “Interpolated finite impulse response filter,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 3, pp. 563–570, June 1984.

[10] J. A. Moorer, “General spectral transformations for digital filters,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, no. 5, pp. 1092–1094, October 1981.

[11] D. Rossum, “Making digital filters sound ’analog’,” in *Proceedings of 1992 International Computer Music Conference*, San Jose, California, USA, October 1992, pp. 30–33.

[12] P. Dutilleux and U. Zölzer, “Delays,” in *DAFX – Digital Audio Effects*, U. Zölzer, Ed., chapter 3, pp. 63–74. John Wiley & Sons, Hoboken, New Jersey, USA, 2002.

[13] T. I. Laakso, V. Välimäki, M. Karjalainen, and U. K. Laine, “Splitting the unit delay – tools for fractional delay filter design,” *IEEE Signal Processing Magazine*, vol. 13, no. 1, pp. 30–60, January 1996.

[14] K. Brandenburg, “Evaluation of quality for audio encoding at low bit rates,” in *Proceedings of the AES 82nd Convention*, London, UK, February 1987.

[15] K. Brandenburg and T. Sporer, “’NMR’ and ’Masking Flag’: Evaluation of quality using perceptual criteria,” in *Proceedings of the AES 11th International Conference on Test and Measurement*, Portland, Oregon, USA, May 1992, pp. 169–179.