INTERNATIONAL SYMPOSIUM on MUSICAL and ROOM ACOUSTICS

ISMRA 2016

September 11-13, 2016

La Plata, Buenos Aires, Argentina

Sound production – Sound synthesis:
Paper ISMR2016-48

Eliminating aliasing caused by discontinuities using integrals of the sinc function

Fabián Esqueda[(a)], **Stefan Bilbao**[(b)], **Vesa Välimäki**[(c)]

[(a)]Aalto University, Dept. of Signal Processing and Acoustics, Espoo, Finland, fabian.esqueda@aalto.fi

[(b)]Acoustics and Audio Group, University of Edinburgh, United Kingdom, s.bilbao@ed.ac.uk

[(c)]Aalto University, Dept. of Signal Processing and Acoustics, Espoo, Finland, vesa.valimaki@aalto.fi

Abstract

A study on the limits of bandlimited correction functions used to eliminate aliasing in audio signals with discontinuities is presented. Trivial sampling of signals with discontinuities in their waveform or their derivatives causes high levels of aliasing distortion due to the infinite bandwidth of these discontinuities. Geometrical oscillator waveforms used in subtractive synthesis are a common example of audio signals with these characteristics. However, discontinuities may also be introduced in arbitrary signals during operations such as signal clipping and rectification. Several existing techniques aim to increase the perceived quality of oscillators by attenuating aliasing sufficiently to be inaudible. One family of these techniques consists on using the bandlimited step (BLEP) and ramp (BLAMP) functions to quasi-bandlimit discontinuities. Recent work on antialiasing clipped audio signals has demonstrated the suitability of the BLAMP method in this context. This work evaluates the performance of the BLEP, BLAMP, and integrated BLAMP functions by testing whether they can be used to fully bandlimit aliased signals. Of particular interest are cases where discontinuities appear past the first derivative of a signal, like in hard clipping. These cases require more than one correction function to be applied at every discontinuity. Results obtained show that if sufficiently many samples are corrected at each discontinuity, aliasing can be virtually eliminated while preserving the spectral envelope of the signal. This work extends the understanding of the BLEP, BLAMP, and integrated BLAMP functions as antialiasing tools.

Keywords: acoustic signal processing, antialiasing, subtractive synthesis

INTERNATIONAL SYMPOSIUM on MUSICAL and ROOM ACOUSTICS
ISMRA 2016
September 11-13, 2016
La Plata, Buenos Aires, Argentina

Eliminating aliasing caused by discontinuities using integrals of the sinc function

1 Introduction

Voltage-controlled oscillators that generate harmonically-rich geometrical waveforms form part of the basic building blocks of subtractive synthesis. Trivial digital implementations of these audio oscillators are prone to aliasing distortion due to the inherent discontinuities found in these waveforms and/or in their derivatives [1, 2, 3, 4, 5]. These discontinuities have infinite bandwidth and must to be sampled at exceedingly high sampling rates in order to keep aliasing down to a negligible level [6]. Aliasing can affect the overall sound quality of digital musical systems, causing audible disturbances such as inharmonicity and beating [4].

The issue of aliasing caused by discontinuities is not restricted to audio synthesis applications. Operations such as signal clipping and rectification also introduce discontinuities in the derivatives of a signal. Previously, tools like oversampling and the harmonic mixer were the only available methods to avoid the aliasing caused by these operations [7, 8]. However, these approaches are far from ideal due to the infinite bandwidth required by these discontinuities. Recent work on antialiasing clipped and rectified audio signals has focused on the use of correction functions commonly used to bandlimit digital oscillators [9, 10]. By correcting a few points at every discontinuity, aliasing can be attenuated. This work examines the use of these correction functions in both synthesis and processing applications in order to demonstrate that, if applied at sufficiently many points, they can be used to fully eliminate aliasing. Of particular interest to this study are cases where the use of more than one correction function is necessary, e.g. cases where discontinuities appear in the first and higher derivatives.

Synthesizing oscillator waveforms with reduced aliasing is a well documented subject in computer music. In 1996, Stilson and Smith suggested generating these waveforms by integrating a bandlimited impulse train, or BLIT [1, 11]. This method was further refined in 2001 by Brandt, who suggested performing the integration process analytically and replacing signal discontinuities with what he called the bandlimited step (BLEP) function [12]. This idea was later optimized using polynomial approximations of the BLEP correction function (or polyBLEP) [2, 4], and extended to synthesize triangular waveforms using a bandlimited ramp (BLAMP) function [4, 13]. Moreover, recent work demonstrated the BLAMP function and its polynomial form (the polyBLAMP) can also be used to reduce the aliasing caused by signal clipping and signal rectification [9, 10]. Other available techniques to synthesize oscillator waveforms with reduced aliasing include the differentiated parabolic waveform (DPW) [3], and polynomial transition regions (PTR) methods [14].

This paper is organized as follows. Section 2 derives the BLEP, BLAMP, and integrated BLAMP correction functions used in this study. Section 3 presents three synthesis examples where the BLEP and BLAMP functions can be used to fully bandlimit trivially-sampled aliased signals. Section 4 then considers the case of hard clipping, which introduces discontinuities in the first higher derivatives of a signal. Finally, concluding remarks are provided in Section 5.

2 Bandlimited correction functions

To bandlimit signals with discontinuities, we begin by modeling a discontinuity in the continuous-time domain using the Heaviside step function, defined as

$$u(t) = \begin{cases} 0 & t < 0 \\ 1 & t \geq 0, \end{cases} \tag{1}$$

where t is time. This function simply "jumps" from 0 to 1 when $t = 0$. The derivative of (1) w.r.t. time is defined as the Dirac delta function [15], so that

$$\frac{\mathrm{d}u(t)}{\mathrm{d}t} = \delta(t). \tag{2}$$

Since the delta function has a flat unity spectrum, its bandlimited form can be derived from the inverse Fourier transform of an ideal brickwall lowpass filter with cutoff at the Nyquist limit (i.e. at half the sampling rate). The resulting expression is known as the bandlimited impulse function and forms the basis of the BLIT synthesis method [1]. It can be written as

$$h^{(0)}(t) = f_{\mathrm{s}}\mathrm{sinc}(f_{\mathrm{s}}t), \tag{3}$$

where $\mathrm{sinc}(x) = \sin(\pi x)/(\pi x)$ is the cardinal sine function and f_{s} is the sampling frequency that will determine the bandwidth of the impulse. Evaluating the integral of (3) yields the closed-form expression for the BLEP function [12], defined as

$$h^{(1)}(t) = \frac{1}{2} + \frac{1}{\pi}\mathrm{Si}(\pi f_{\mathrm{s}}t), \tag{4}$$

where $\mathrm{Si}(x)$ is the sine integral $\mathrm{Si}(x) = \int_0^x \frac{\sin(t)}{t}\mathrm{d}t$ and the superscript $^{(1)}$ denotes it is the first integral of (3). Figure 1(a) shows the time-domain form of this expression. For the sake of simplicity, a unit sampling interval has been used to produce this figure. Figure 1(d) shows the difference between (4) and (1), or BLEP residual function. To bandlimit a discrete-time signal with a discontinuity, this residual function has to be centered around the exact point where the discontinuity occurs, sampled at the nearest integer sample points and added to the signal. A detailed explanation on how to use this residual function can be found in [4].

By integrating (4) we can derive expressions for the BLAMP function [4, 9, 10], given by

$$h^{(2)}(t) = t\left[\frac{1}{2} + \frac{1}{\pi}\mathrm{Si}(\pi f_{\mathrm{s}}t)\right] + \frac{\cos(\pi f_{\mathrm{s}}t)}{\pi^2 f_{\mathrm{s}}}, \tag{5}$$

and the integrated BLAMP function [16]:

$$h^{(3)}(t) = \frac{t^2}{2}\left[\frac{1}{2} + \frac{1}{\pi}\mathrm{Si}(\pi f_{\mathrm{s}}t)\right] + t\frac{\cos(\pi f_{\mathrm{s}}t)}{2\pi^2 f_{\mathrm{s}}} + \frac{\sin(\pi f_{\mathrm{s}}t)}{2\pi^3 f_{\mathrm{s}}^2}. \tag{6}$$

The BLAMP and integrated BLAMP functions can be used to bandlimit discontinuities that appear in the first and second derivative of a signal, respectively. Figures 1(b), (c), (e), and

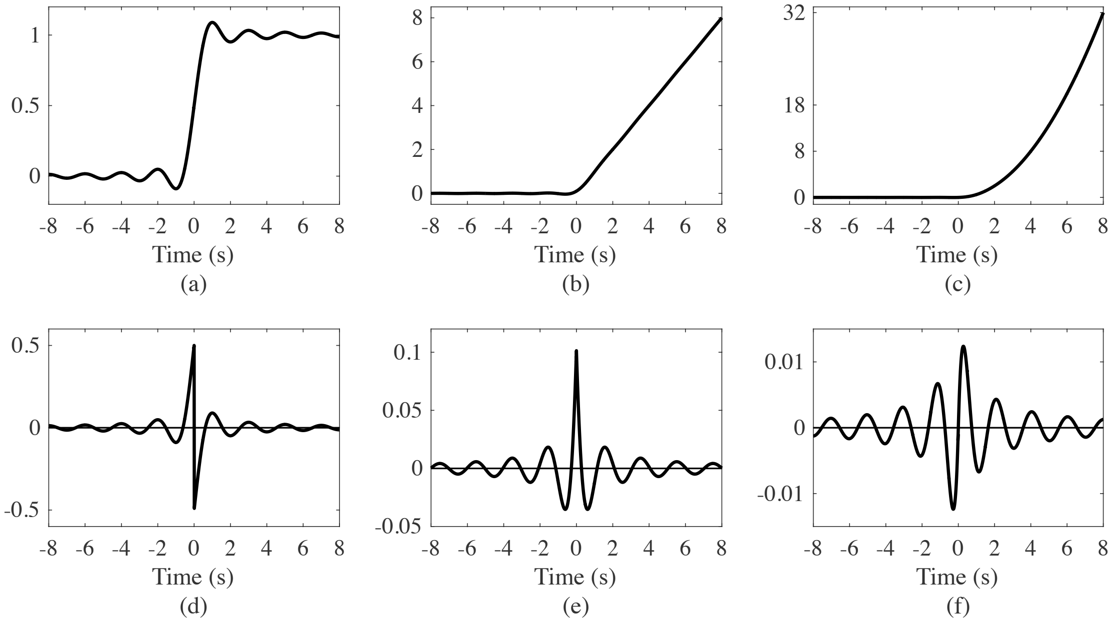

Figure 1: **Time domain representation of the (a) BLEP function, (b) BLAMP function, (c) integrated BLAMP function, and (d)–(f) their corresponding residuals.**

(f) show the time-domain representation of the BLAMP and integrated BLAMP functions along with their residual functions, computed by subtracting the trivial ramp [i.e. the integral of (1)] and the integral of the trivial ramp, respectively.

As suggested by Figures 1(d), (e) and (f), the three residual functions do not have finite support. This means that truncating them and superimposing them at every discontinuity will introduce new discontinuities, causing further artifacts. Additionally, these three functions are computationally expensive due to the presence of the sine integral function. Both issues can be addressed at once by storing a finite portion of these functions in a lookup table and applying a window function. The symmetry properties of all three residual functions can be exploited to reduce table size. A thorough study on the table-based BLEP method using different window functions can be found in [17].

The examples presented in the following sections of this study were implemented using table-based residual functions oversampled by a factor 100 and windowed using a Hann window. Additionally, cubic Lagrangian interpolation was used to reduce any round-off errors. A threshold of $-160\,$dB was chosen to evaluate the presence of aliasing, i.e. signals without aliasing components above this threshold are considered bandlimited.

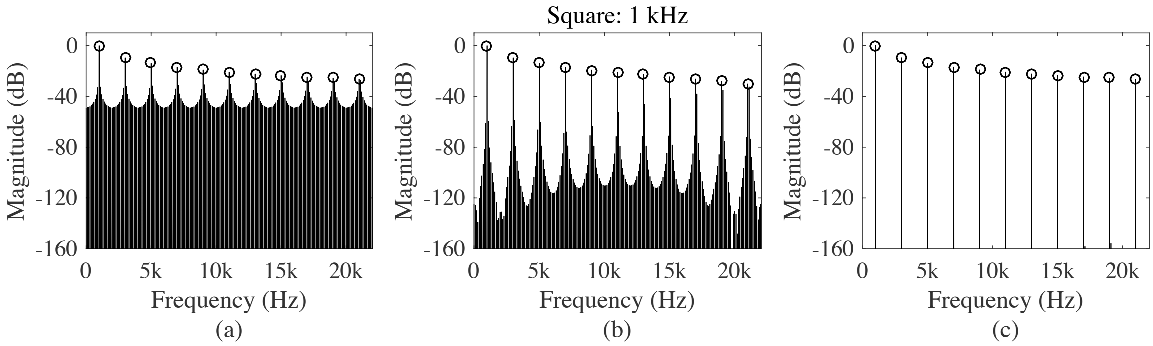

Figure 2: **Magnitude spectrum of a 1-kHz square signal (a) trivially sampled, (b) with 4-point BLEP correction, and (c) with 2000-point BLEP correction. Circles indicate non-aliased harmonics.**

3 Bandlimited oscillator waveforms

To exemplify how the BLEP and BLAMP functions can be used to fully bandlimit audio signals with discontinuities we begin by considering two simple test signals: the rectangular (or square wave) and triangular oscillators. These two waveforms are not infinitely differentiable, and as such, have a finite number of discontinuities.

Figure 2(a) shows the magnitude spectrum of a trivially-sampled square wave with a fundamental frequency of 1 kHz[1]. This waveform contains odd harmonics only and has two discontinuities per period, which cause extremely high levels of aliasing distortion. To suppress this aliasing using the BLEP method, the BLEP residual function must be centered at each discontinuity, scaled by the peak-to-peak amplitude of the signal, sampled at the nearest sample points, and added to the trivial waveform. Figure 2(b) shows the magnitude spectrum of the square wave after 4-point BLEP correction. By correcting only 2 points on each side of every discontinuity, a significant portion of the aliased frequencies has already been suppressed, with some spurious components near DC and Nyquist attenuated by more than 100 dB.

Depending on the application, the 4-point correction may be considered sufficient. However, the aim of this work is to demonstrate that these bandlimited correction functions can eliminate all aliases if sufficiently many points are corrected. Increasing the scope of the correction to more points gradually converges towards a practically bandlimited signal, with no aliases visible above the -160-dB line. This is presented in Figure 2(c), where two thousand points have been corrected at every discontinuity. When using long table-based BLEP residuals, correction functions at each discontinuity will overlap with each other.

In triangular oscillators, discontinuities appear in the first derivative of the signal. These discontinuities can be easily detected by identifying any sharp edges or corners [10]. Figure 3(a) shows the magnitude spectrum of a 3.9-kHz trivially-sampled triangular waveform. Even at a higher fundamental, this signal exhibits less aliasing than the square wave due to the steep rolloff of the integrated discontinuities, hence its *smooth* sound quality. The trivial triangular

[1] A fixed sampling rate $f_s = 44.1$ kHz was used in this and subsequent examples in this study.

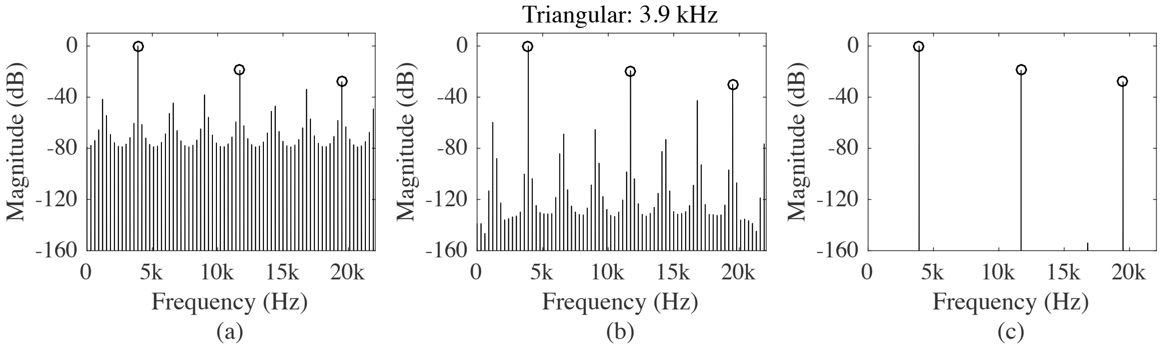

Figure 3: **Magnitude spectrum of a 3.9-kHz triangular signal (a) trivially sampled, (b) with 4-point BLAMP correction and (c) with 2000-point BLAMP correction.**

signal can be antialiased by sampling and adding the BLAMP residual function at every corner. As before, correcting 4 samples at every corner suppresses most of the aliasing [see Figure 3(b)], while pushing the remaining aliases below -160 dB also requires correcting two thousand samples per discontinuity [see Figure 3(c)].

A more interesting behavior can be observed in audio signals where discontinuities appear in both the waveform and its first derivative. An example signal with these characteristics is the sawtooth-triangular waveform, included as one of the oscillator waveforms available in the iconic Minimoog Model D synthesizer [18]. As hinted by its name, the sound of this waveform can be described as a mixture between the sawtooth and the triangular signals. Figure 4(a) shows two periods of a continuous-time sawtooth-triangular signal with an arbitrary fundamental frequency $f_0 = 1/T_0$. As shown in this diagram, this waveform has one discontinuity per period. Computing the first derivative of this signal w.r.t. time yields a rectangular signal which has two discontinuities per period [see Figure 4(b)]. For illustrative purposes only, normalized arrows have been used to help indicate the location, magnitude and polarity of the differentiated discontinuities. Parameters μ_1 and μ_2 represent the slope values of the rising and falling portions of the original signal, which are determined by the fundamental frequency. Further differentiation of this signal results in the alternating impulse train shown in Figure 4(c).

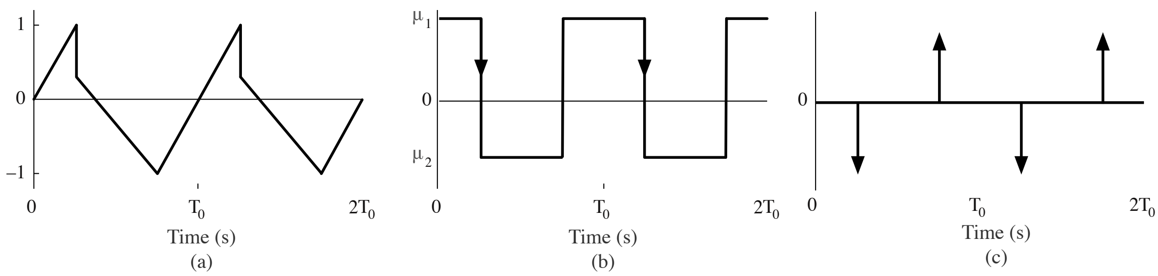

Figure 4: **Two periods of a (a) sawtooth-triangular signal, and (b)–(c) its first and second derivatives, respectively.**

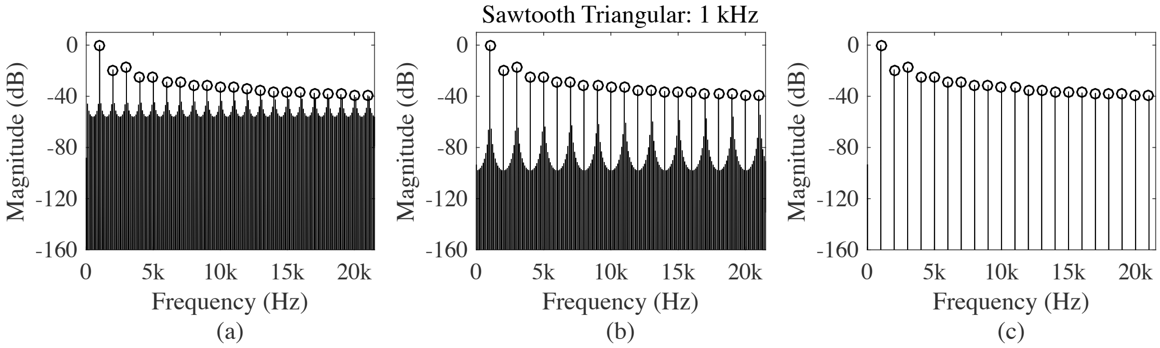

Figure 5: **Magnitude spectrum of a 1-kHz sawtooth-triangular signal (a) trivially sampled, (b) with 1200-point BLEP correction and (c) with 1200-point BLEP and BLAMP correction.**

Figure 5(a) shows the magnitude spectrum of a heavily-aliased 1-kHz sawtooth-triangular signal trivially sampled. Figure 5(b) presents the magnitude spectrum of the same signal after 1200-point BLEP correction. A significant portion of aliasing has been removed and the *harsh* sound caused by the aliases has already been reduced. The remaining aliases are solely caused by the discontinuities in the first derivative. Applying an additional layer of correction using the BLAMP function successfully attenuates the remaining aliases [see Figure 5(c)], restoring the original analog sound of the signal. This example demonstrates the independence of aliases caused by discontinuities in the waveform from those caused by discontinuities in the first derivative, and how they must be handled separately when using bandlimited correction functions.

4 Bandlimited hard clipping

In audio processing applications, where discontinuities are introduced by operations such as signal clipping and signal rectification, discontinuities will extend to virtually all derivatives of the signal. As with the sawtooth-triangular waveform, these discontinuities must be dealt with separately in order to fully eliminate aliasing from the signal.

To observe this behavior, we consider the elementary case of a sinewave processed by a bipolar hard clipper. In bipolar clipping, sample values above and below a predetermined value are set to this value while the rest of the samples are left unaffected. This generates harmonic distortion and affects the timbre of the signal. Assuming an input sinewave normalized between $[-1,1]$ we define a clipping threshold $L \in (0,1]$, where $L = 1$ means no clipping. Figure 6(a) shows the magnitude spectrum of a 500-Hz clipped sinewave with clipping threshold $L = 0.9$. This threshold means that only the tips of the signal have been clipped. Hard clipping generates corners in the waveform, which translate intro discontinuities in the first derivative of the signal. Following our previous examples, we can bandlimit these corners using the BLAMP residual function. In this application, the BLAMP residual must be scaled by the slope of the original unclipped signal at the clipping points [9]. This parameter determines the magnitude of the discontinuity introduced. Figure 6(b) shows the spectrum of the clipped sinusoid after 200-

point BLAMP correction. On average, the entire level of aliasing has been reduced by roughly 30 dB. Here, increasing the correction scope has little to no effect on the magnitude spectrum of the signal.

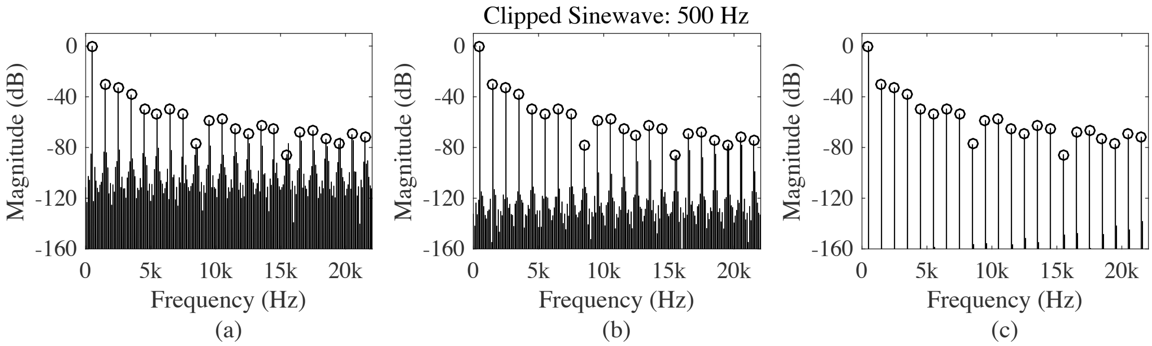

Figure 6: **Magnitude spectrum of a 500-Hz clipped sinewave (a) trivially sampled, (b) after 200-point BLAMP correction, and (c) after 200-point BLAMP and integrated BLAMP correction.**

A second layer of correction can then be implemented using the integrated BLAMP residual function [see Figure 1(f)], which upon sampling has to be scaled by the second derivative of the signal (i.e. the curvature) at the clipping points. Figure 6(c) shows the result of this process, where the remaining aliases have been further attenuated. The fact that the overall spectral envelope of the signal was preserved after the correction process means that by using the BLAMP and integrated BLAMP functions we are actually approximating the ideal solution. In the time domain, this would be equivalent to recreating Gibbs phenomenon [9, 15].

From a perceptual point of view, the results obtained in the previous example after using only the BLAMP function may be sufficient to label the signal as being "alias-free" since most aliases lie below -100 dB. This is not the case at higher frequencies, where the value of both the slope and curvature of the signal also increase. As a final example, Figure 7(a) shows the magnitude spectrum of a 4186-kHz sinusoid (highest fundamental frequency on the piano) with the same clipping threshold $L = 0.9$. The resulting signal only has three harmonics below the Nyquist limit and high levels of aliasing above -100 dB. Applying the BLAMP and integrated BLAMP residual functions at every clipping point yields the spectra seen in Figures 7(b) and (c). Once again, dealing with both layers of discontinuities has significantly improved signal quality. For instance, the level of the highest alias below the fundamental, at 1948 kHz, was attenuated from -56 to -103 dB, a difference of 47 dB. Extending the range of the correction process did not improve signal quality any further. This could be attributed to further discontinuities in higher derivatives or to estimation errors introduced during the lookup table reading process.

These two examples showcase how bandlimited correction functions can work together to eliminate aliasing caused by discontinuities introduced during signal clipping. For low clipping thresholds, where the curvature of the signal may be close to zero, the integrated BLAMP correction will simply cancel out. In practical implementations, where processing thousands of samples per second is simply not feasible, polynomial approximations of these correction functions can be used instead.

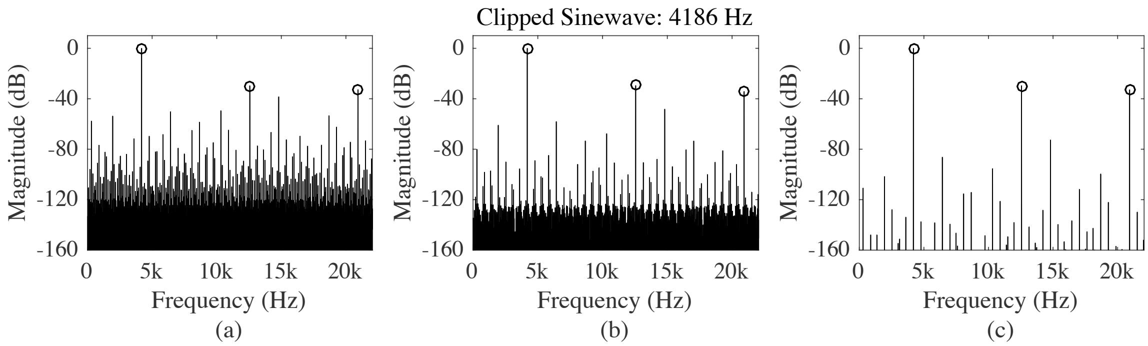

Figure 7: **Spectrum of a 4186-Hz clipped sinewave (a) trivially sampled, (b) after 200-point BLAMP correction, and (c) after 200-point BLAMP and integrated BLAMP correction.**

5 Conclusions

This work studied the use of various bandlimited correction functions to eliminate aliasing caused by discontinuities in audio signals. In synthesis applications, these discontinuities may be inherent to the shape of the waveform being synthesized. In a processing scenario, discontinuities may be introduced by operations such as signal clipping. The BLEP, BLAMP, and integrated BLAMP functions can be used to bandlimit discontinuities that appear in the zeroth, first, and second derivative of a signal, respectively. Audio signals with discontinuities in more than one derivative, such as the sawtooth-triangular oscillator and hard-clipped signals, require the use of more than one of these correction functions. For instance, eliminating aliasing from a clipped signal requires the use of both the BLAMP and integrated BLAMP functions. Results obtained from this study demonstrate that if sufficiently many samples are corrected at each discontinuity aliasing can be completely eliminated.

Acknowledgments

This paper was written during the first author's research visit to the University of Edinburgh in summer 2016. This visit was supported by the Foundation for Aalto University Science and Technology and the Aalto ELEC Doctoral School. S. Bilbao was supported by the European Research Council, under grant number ERC-StG-2011-279068-NESS.

References

[1] T. Stilson and J. Smith. Alias-free digital synthesis of classic analog waveforms. In *Proc. International Computer Music Conf.*, pages 332–335, Hong Kong, 1996.

[2] V. Välimäki and A. Huovilainen. Antialiasing oscillators in subtractive synthesis. *IEEE Signal Process. Mag.*, 24(2):116–125, Mar. 2007.

[3] V. Välimäki, J. Nam, J. O. Smith, and J. S. Abel. Alias-suppressed oscillators based on differentiated polynomial waveforms. *IEEE Trans. Audio Speech Lang. Process.*, 18(4):786–798, May 2010.

[4] V. Välimäki, J. Pekonen, and J. Nam. Perceptually informed synthesis of bandlimited classical waveforms using integrated polynomial interpolation. *J. Acoust. Soc. Am.*, 131(1):974–986, Jan. 2012.

[5] H.-M. Lehtonen, J. Pekonen, and V. Välimäki. Audibility of aliasing distortion in sawtooth signals and its implications for oscillator algorithm design. *J. Acoust. Soc. Am.*, 132(4):2721–2733, Oct. 2012.

[6] J. Schimmel. Audible aliasing distortion in digital audio synthesis. *Radioengineering*, 21(1):57, 2012.

[7] H. Thornburg. Antialiasing for nonlinearities: acoustic modeling and synthesis applications. In *Proc. Int. Comput. Music Conf.*, pages 66–69, Beijing, China, Oct. 1999.

[8] J. Schattschneider and U. Zölzer. Discrete-time models for nonlinear audio systems. In *Proc. DAFX-99 Digital Audio Effects Workshop*, pages 45–48, Trondheim, Norway, 1999.

[9] F. Esqueda, S. Bilbao, and V. Välimäki. Aliasing reduction in clipped signals. *IEEE Trans. Signal Process.*, 2016. accepted for publication.

[10] F. Esqueda, V. Välimäki, and S. Bilbao. Antialiased soft clipping using an integrated bandlimited ramp. In *Proc. European Signal Processing Conf. (EUSIPCO 2016)*, Brno, Czech Republic, 2016.

[11] J. Nam, V. Välimäki, J. S. Abel, and J. O. Smith. Efficient antialiasing oscillator algorithms using low-order fractional delay filters. *IEEE Trans. Audio Speech Lang. Process.*, 18(4):773–785, May 2010.

[12] E. Brandt. Hard sync without aliasing. In *Proc. Int. Computer Music Conf.*, pages 365–368, Havana, Cuba, Sept. 2001.

[13] A. Huovilainen. Design of a scalable polyphony-MIDI synthesizer for a low cost DSP. Master's thesis, Aalto University, Espoo, Finland, 2010.

[14] J. Kleimola and V. Välimäki. Reducing aliasing from synthetic audio signals using polynomial transition regions. *IEEE Signal Process. Lett.*, 19(2):67–70, Feb. 2012.

[15] R. N. Bracewell. *The Fourier Transform and its Applications*. McGraw-Hill, 2nd edition, 1986.

[16] F. Esqueda, V. Välimäki, and S. Bilbao. Rounding corners with BLAMP. In *Proc. 19th Int. Conf. Digital Audio Effects (DAFx-16)*, Budapest, Hungary, Sept. 2016.

[17] Will Pirkle. *Designing Software Synthesizer Plug-Ins in C++*. Focal Press, 1st edition, Nov. 2014.

[18] Moog Music Inc. Minimoog Model D. Product information available at `https://www.moogmusic.com/products/minimoog/minimoog-model-d/`, accessed July 15, 2016.