NEW APPROACHES TO DIGITAL SUBTRACTIVE SYNTHESIS

Antti Huovilainen and Vesa Välimäki
Helsinki University of Technology
Laboratory of Acoustics and Audio Signal Processing
P.O. Box 3000, FI-02015 TKK, Espoo, Finland

ABSTRACT

Computationally efficient oscillator and filtering algorithms for digital subtractive synthesis are discussed. The oscillators algorithms include the recently proposed differentiated parabolic waveform generator and its modification. The algorithm generates a signal that sounds similar to the analog sawtooth waveform, because it suppressed aliasing that occurs due to sampling of a non-bandlimited waveform. A modified version of the nonlinear digital Moog ladder filter is introduced. The new structure reduces the computational cost of the nonlinear digital Moog filter by using a single nonlinearity in the feedback loop instead of four nonlinear functions inside filter sections. The new digital Moog filter structure also decouples the cutoff and the resonance parameters and offers several response types by selecting a weighted sum of different output points.

1. INTRODUCTION

Digital subtractive synthesis, which is also called virtual analog synthesis, refers to computational methods that imitate the sound generation principles of analog synthesizers of the 1960s and 1970s. The basic principle in subtractive synthesis is first to generate a signal with a rich spectral content, and then to filter that signal with a time-varying resonant filter.

Virtual analog synthesis became a popular and commercial term in about 1995, when Clavia introduced the Nord Lead 1 synthesizer, which was marketed as an analog-sounding digital synthesizer that uses no samples. Instead, all sounds were generated by simulating analog subtractive synthesis. Previously, the Roland D-50 synthesizer of the late 1980s worked in a similar way although it contained sampled sounds. An early example of an attempt to design a digital synthesizer that sounds analog was Synergy [4].

What makes digital subtractive synthesis more demanding than is generally understood is that imitating analog electronics with digital processing is not as easy as it may seem. One problem is aliasing caused by sampling of analog waveforms that have sharp edges. The spectra of such waveforms continue infinitely high, and the signals are thus not bandlimited. Another difficulty is that analog filters do not obey simple linear theory. With high signal levels they generate distortion. This does not naturally occur in digital processing, but it must be designed and implemented on purpose, see for example, references [8] and [3].

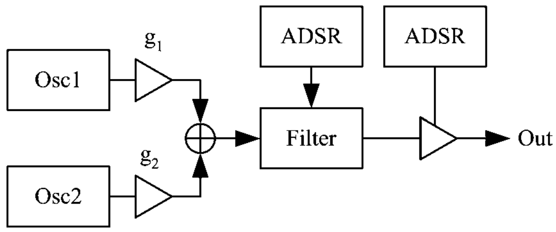

Figure 1. A typical block diagram of subtractive synthesis as it was implemented in the Prophet 5 synthesizer in late 1970s.

In this paper, we discuss new versions of oscillator and resonant filtering algorithms that can sound like old analog synthesizers.

2. SUBTRACTIVE SYNTHESIS

The electronic music modules introduced by Robert A. Moog in mid-1960s [6] are one of the most important innovations in music technology. A few years later, his company introduced products where the various modules, such as oscillators, filters, and amplifiers, were integrated into a single portable unit. Subtractive synthesis was the main principle used in these instruments. Minimoog was one of the most popular analog synthesizers in 1970s.

The Prophet 5 synthesizer introduced by Sequential Circuits in 1979 has microprocessor controlled electronics, but it is still an analog synthesizer. Its block diagram shown in Fig. 1 is today a classic example of the subtractive synthesis principle. It includes two oscillators, a resonant lowpass filter, and two envelope generators (ADSR). There are a couple of alternative waveforms available together with a noise source.

3. DIGITAL OSCILLATORS

The sharp edges of geometric waveforms, such as the sawtooth or the square wave, cause aliasing, because such signals are not bandlimited. Three different classes of methods are known to avoid this problem:

1. Bandlimited methods that generate harmonics only below the Nyquist limit, such as additive synthesis and its variants, e.g., wavetable synthesis and the discrete summation formulae;
2. Quasi-bandlimited methods in which aliasing is low and its level can be adjusted by design to save computational costs, such as in the BLIT [9] and the minBLEP [1] techniques;

3. Alias-suppressing methods in which it is accepted that some aliasing will occur but an attempt is made to attenuate it sufficiently.

In this work, we focus on the third category of methods. Two approaches are currently available: the distortion and filtering of sine waves, which we call Lane's method [5], and the differentiated parabolic waveform (DPW) [11]. We discuss the latter one in the following.

3.1. DPW algorithm

The simplest version of the DPW algorithm [11] generates the sawtooth waveform in four stages, as illustrated in Fig. 2: First generate the trivial sawtooth waveform using a modulo counter, then raise the waveform to the second power, differentiate the signal with a first difference filter with transfer function $H_D(z) = 1 - z^{-1}$, and, finally, scale the obtained waveform by factor $c = f_s/(4f)$, where f is the fundamental frequency of the sawtooth signal and f_s is the sampling rate.

The waveform produced by the modulo counter resembles the sawtooth waveform, as seen in Fig. 3(a), but it sounds badly distorted. The reason is that its spectrum decays slowly, about 6 dB per octave. When it is sampled, the spectral components above the Nyquist limit are mirrored down to the audible frequencies. This is clearly seen in Fig. 4(a), where the desired harmonics are indicated by circles and the rest of the peaks are aliased images.

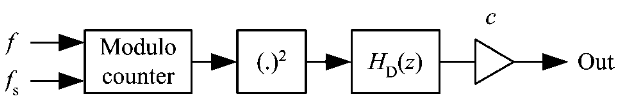

Figure 2. The DPW algorithm [11].

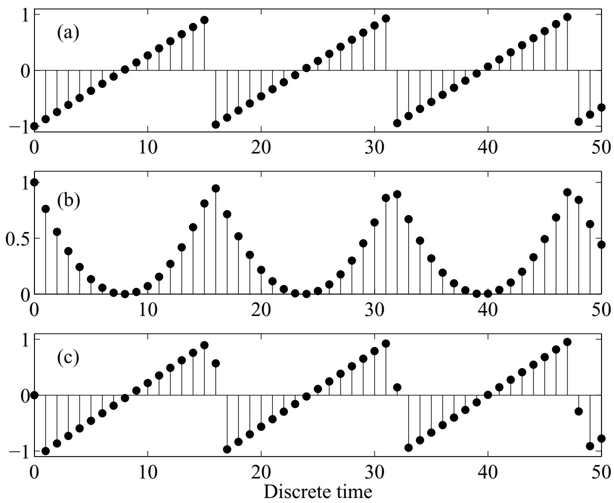

Figure 3. (a) The trivial sawtooth waveform, (b) the squared sawtooth wave, and (c) the differentiated parabolic waveform. The fundamental frequency is 2793.8 Hz (MIDI note number 101), and the sampling rate is 44.1 kHz.

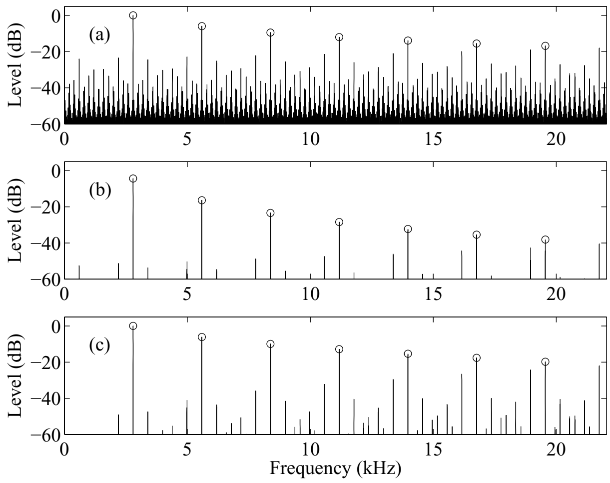

Figure 4. The spectra of the waveforms shown in Fig. 3: (a) the trivial sawtooth waveform, (b) the squared sawtooth wave, and (c) the differentiated parabolic waveform. The desired spectral components (2793.8 Hz, 5587.6 Hz, 8381 Hz, …) are circled (o), while the rest of the spectral components are caused by aliasing and are heard as disturbance.

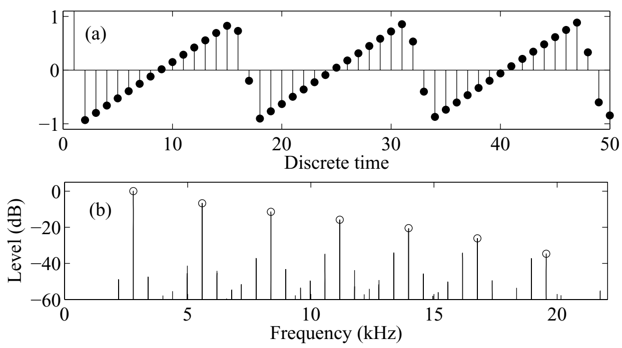

Figure 5. (a) The waveform and (b) the spectrum of the signal obtained with the averaged differentiator.

Raising the signal to the second power modifies the waveform so that it now consists of pieces of parabola, see Fig. 3(b). The spectrum of this waveform decays about 12 dB per octave, and this is why aliasing is suppressed in Fig. 4(b) [11]. Finally, when the piecewise parabolic signal is differentiated and scaled, the signal again looks like the sawtooth waveform, see Fig. 3(b), but the aliased components are suppressed, as seen in Fig. 4(c).

A remaining problem is that at high frequencies the level of aliased components is close to that of the harmonics. This may lead in some cases to beating. A solution of avoid this is to replace the differentiator with its averaged version $H_D(z) = 1 - z^{-2} = (1 + z^{-1})(1 - z^{-1})$. The resulting waveform and spectrum are shown in Fig. 5. It is seen that in the discrete-time waveform in Fig. 5(a) the transitions from the maximum value (near +1) to the minimum value (near –1) are smoother than in Fig. 3(c). The corresponding spectrum, see Fig. 5(b), decays faster at high frequencies than that in Fig. 4(c).

4. DIGITAL RESONANT FILTERS

A musical filter differs from traditional IIR filters in mainly three ways: the parameters are changed at a rapid rate, the order is usually predetermined instead of matching to certain stopband attenuation specification and a controllable resonant peak is introduced near the cutoff frequency.

A ”perfect” digital resonant filter then fulfills the following criteria:

1. Coefficient update should be fast. To avoid clicks, the coefficients should be updated on a per-sample basis [7].
2. The filter cutoff and resonance parameters should be decoupled. Change in one should not affect the other.
3. The filter should stay unconditionally stable as long as parameters are inside the allowed range.
4. The filter should have a response similar to an existing analog resonant filter. Some analog filters have a characteristic sound that should be emulated, if possible.
5. The filter should be capable of self-oscillation.

A number of filters trying to fill the criteria have been developed. We take a closer look at the Moog lowpass filter [6] in the following.

4.1. Digital Moog filter

The Moog ladder filter [6] can be considered the first musical filter. It features independent voltage control of both the cutoff frequency and the resonance amount, while also having a characteristic sound of its own. The filter consists of four identical one-pole lowpass sections (implemented with an innovative transistor ladder circuit) in series with a global negative feedback to produce a resonant peak near the cutoff frequency.

A digital model of the Moog filter was first presented by Stilson and Smith [10]. As in the analog prototype, it has four one-pole filters in series, and a global feedback is used to produce the resonance. To realize the filter, a unit delay has to be inserted in the feedback path, but this couples the cutoff and the resonance controls. Various ways of compensation have been examined, with the “compromise” version [10] being the most attractive. The “compromise” version inserts a zero at $z = -0.3$ inside each one-pole section, thus mostly decoupling the resonance and the cutoff parameters. The modified one-pole structure is shown in figure 6.

The coefficient g determines the cutoff frequency and is approximately $g = 1 - \exp(-2\pi f_c/f_s)$, where f_c is the desired cutoff frequency and f_s is the sampling rate. If exact tuning is desired, a tuning table can be incorporated. Practical implementations usually use an interpolated table lookup to evaluate g in any case.

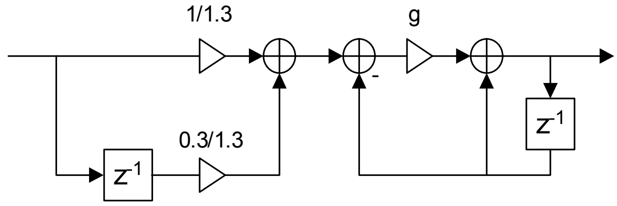

Figure 6. Compromise one-pole section.

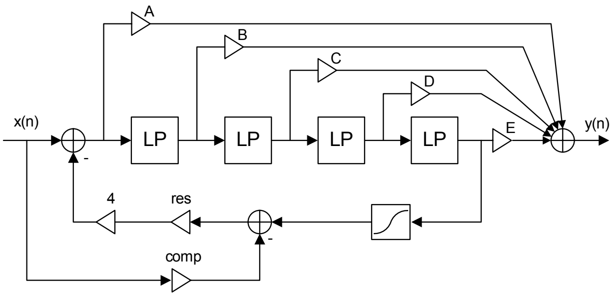

Figure 7. Improved Moog-style filter. Each LP block contains the one-pole structure shown in Fig. 6. Coefficients A, B, C, D, and E can be used to select the type of output (lowpass, highpass, bandpass, or notch filter or one of their combinations).

4.2. Improved Moog filter

While the Stilson and Smith Moog model is certainly useful and solves the problem of fast coefficient update, it becomes unstable with very large resonance values and it cannot self-oscillate. Furthermore, it does not emulate the characteristic distortion produced by the original transistor ladder circuit.

Huovilainen has developed an improved model that models the ladder circuit by inserting nonlinearities inside the one-pole sections [3]. This improved model more closely emulates the characteristic sound and is also capable of self-oscillation. A disadvantage is the need for five hyperbolic tangent (tanh) function evaluations per sample and oversampling by factor two at least.

An alternative and extended model is shown in figure 7. The embedded nonlinearities within sections are replaced by a single nonlinearity in the feedback loop, thus greatly reducing the computational cost of the filter. We have used the tanh function for the nonlinearity, but any smoothly saturating function may be used. There is a difference in the sound compared to the full nonlinear Moog filter model, but this model can emulate most of the behavior, such as self-oscillation. Its output is also always bounded.

The new model also contains two additional improvements. The traditional Moog filter and similar cascaded one-pole filters suffer from decreasing the pass-band gain as the resonance is increased, because the resonance is produced with a global negative feedback. If some of the input is subtracted from the feedback signal before multiplying by the resonance amount, the pass-band gain change can be controlled [2]. A value of

1.0 for the *comp* parameter keeps the pass-band gain constant. This, however, results in a large increase of the output amplitude as the resonance is increased. To keep the overall level approximately constant, *comp* should be set to 0.5 resulting in a 6 dB passband gain decrease at the maximum resonance (compared to a 12 dB decrease in the original Moog model).

Another improvement is the addition of various frequency response modes besides the original 24 dB/oct lowpass mode. This can be easily achieved by mixing the individual section outputs with different weights. The concept was pioneered in the Oberheim Xpander and Matrix-12 synthesizers [7], but it was not widely used due to a large number of required components and the need for precision resistors. Accurate mixing is trivial in the digital domain, and a large number of different low-pass, band-pass, high-pass, and notch filter responses and their combinations can be realized. The response can be morphed between these modes by changing the coefficients at runtime, thus allowing interesting modulation possibilities.

5. CONCLUSIONS

Digital subtractive synthesis is a modern approach in which components of analog music synthesizers are modeled using signal processing methods. In this paper, we discussed new oscillator and resonant filtering algorithms. The DPW oscillator algorithm generates a sawtooth waveform approximation that has reduced aliasing with respect to the trivial sawtooth waveform (i.e., the modulo counter output). This new method is probably the simplest useful technique for this purpose, because only the trivial sawtooth is simpler, but it is practically useless due to its heavy aliasing. In this paper, an alternative form for the differentiator to be used in the DPW algorithm was proposed.

The new nonlinear model of the Moog ladder filter is based on a cascade of four one-pole filters and a feedback loop that contains a memoryless nonlinearity. The proposed new Moog filter structure has nice advantages, such as a smaller computational cost than that of a recently proposed nonlinear filter structure, the decoupling of the cutoff frequency and the resonance parameters, and the possibility to obtain various types of filter responses by selecting a weighted sum of different output points. The proposed methods allow the synthesis of retro sounds with modern signal processing techniques.

6. ACKNOWLEDGMENTS

This work has been financed partly by the Academy of Finland (project no. 104934). The authors are grateful to Teppo Karema and Tomi Valkonen of VLSI Solution Oy (Tampere, Finland) for helpful discussions.

7. REFERENCES

[1] Brandt, E. “Hard sync without aliasing,” *Proceedings of the International Computer Music Conference*, Havana, Cuba, 2001. http://www-2.cs.cmu.edu/~eli/L/icmc01/hardsync.html

[2] Curtis Electromusic Specialties. “CEM 3328 four pole low-pass VCF,” 1984. http://www.synthtech.com/cem/c3328pdf.pdf

[3] Huovilainen, A. “Nonlinear digital implementation of the Moog ladder filter,” *Proceedings of the International Conference on Digital Audio Effects*, pp. 61–64, Naples, Italy, 2004. http://dafx04.na.infn.it/

[4] Kaplan, S. J. “Developing a commercial digital sound synthesizer,” *Computer Music Journal* 5(3): 62–73, 1981.

[5] Lane, J., D. Hoory, E. Martinez, and P. Wang. “Modeling analog synthesis with DSPs,” *Computer Music Journal* 21(4): 23–41, 1997.

[6] Moog, R. A. “Voltage-controlled electronic music modules,” *Journal of the Audio Engineering Society* 13(3): 200–206, 1965.

[7] Oberheim. “Xpander service manual,” 1984.

[8] Rossum, D. “Making digital filters sound analog,” *Proceedings of the International Computer Music Conference*, San Jose, CA, pp. 30–33, 1992.

[9] Stilson, T., and J. O. Smith. “Alias-free digital synthesis of classic analog waveforms,” *Proceedings of the International Computer Music Conference*, Hong Kong, China, pp. 332–335, 1996. http://ccrma.stanford.edu/~stilti/papers/

[10] Stilson, T., and J. O. Smith. “Analyzing the Moog VCF with considerations for digital implementation,” *Proceedings of the International Computer Music Conference*, Hong Kong, China, pp. 398–401, 1996. http://ccrma.stanford.edu/~stilti/papers/

[11] Välimäki, V. “Discrete-time synthesis of the sawtooth waveform with reduced aliasing,” *IEEE Signal Processing Letters* 12(3): 214–217, 2005.