Filters, Delays, Modulations and Demodulations
A Tutorial

Pierre Dutilleux
ZKM | Institute for Music and Acoustics
pdx@zkm.de

Abstract

A set of basic signal processing building blocks is presented and the implementation issues, that are relevant to musical applications, are underlined. in this text a filter classification is first outlined then implementation schemes for analog and digital lowpass-filters, within the canonical and the state-variable structure, are given. As well as higher order filters according to the Butterworth design are suggested, delay-based effects such as comb-filter, resonator, flanger, chorus, slapback and echo are reviewed. Phasing is mentioned and normalization schemes for some filters are proposed. Modulators (ring, amplitude, single side band) as well as demodulators (AM-detector, amplitude follower with independent attack and release times, instantaneous envelope follower) are reviewed and their applications listed.

1 Introduction

The goal of this tutorial is to review some widespread digital audio effects but also to give ideas how to improve them and how to extend their range of application. It will attempt to show how each effect relates to physical and acoustical phenomena or which forefather it has in the analog domain. After a brief explanation of the effect, an implementation scheme will be proposed and reference to sound effects or to musical applications will be made. Whether an effect is stiff or easy to tune is relevant here, that is why this issue will be emphasized, where possible.

The first part will deal with filters. After a general introduction and a classification, we will focus on the second order filters. Various implementation schemes will be presented, especially the state variable filter structure. The second part will review delay-based audio effects after introducing the basic building blocks that are the FIR-comb filter and the IIR-comb filter. The third part will introduce modulations, as far as they can be of use for digital audio effects. In the last part we will consider the operation that is reciprocal to the modulation: the demodulation. We will then give details about the envelope follower and its applications.

2 Filters

2.1 A filter classification

The various types of filters can be defined according to the following classification.

- Lowpass (LP): Selects low frequencies up to the cutoff frequency f_c. Attenuates frequencies higher than f_c.
- Highpass (HP): Selects frequencies higher than f_c and attenuates frequencies below f_c.
- Bandpass (BP): Selects frequencies between a lower cutoff frequency f_{cl} and a higher cutoff frequency f_{ch}. Attenuates frequencies below f_{cl} and frequencies higher than f_{ch}.
- Bandreject (BR): Attenuates frequencies between a lower cutoff frequency f_{cl} and a higher cutoff frequency f_{ch}. Selects frequencies below f_{cl} and frequencies higher than f_{ch}.
- Notch: Attenuates frequencies in a narrow bandwidth around the cutoff frequency f_c.
- Resonator: Amplifies frequencies in a narrow bandwidth around the cutoff frequency f_c.

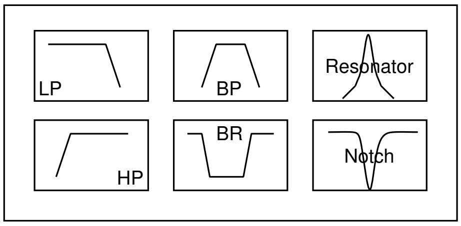

Figure 1: A filter classification.

Other types of filters (LP with resonance, comb, multiple notch...) can be described as a combination of these basic elements. Here are listed some of the possible applications of these filter types: The lowpass with resonance is very often used in computer music to simulate an acoustical resonating structure; the highpass filter can remove undesired very low frequencies; the bandpass can produce effects such as the imitation of a telephone line or of a mute on an acoustical instrument; the bandreject can divide the audible spectrum into two bands that seem to be non correlated. The resonator can be used to add artificial resonances to a sound; the notch is most useful to eliminate annoying frequency components; a set of notch filters, used in combination with the input signal, can produce a phasing effect.

2.2 Filter implementations

A filter can be implemented in various ways. It can be an acoustic filter, as in the case of the voice. For our applications we will rather use electronic or digital means. Although we are interested in digital audio effects, it is worth it to have a look at well established analog techniques because a large body of methods have been developed in the past to design and build analog filters. There are intrinsic design methods for digital filters but many structures can be adapted from existing analog designs. Furthermore, some of them have been tailored for ease of operation within musical applications. It is therefore of interest to gain ideas from these analog designs in order to build digital filters having similar advantages. We will focus on the second order lowpass filter because it is the most common type and other types can be derived from it. The tuning parameters of this lowpass filter are the cutoff frequency f_C and the damping factor ζ. (The lower the damping factor the higher the resonance at the cutoff frequency.)

2.2.1 Analog Design, Sallen & Key

Let us remind of an analog structure that implements a second order lowpass filter with the least number of components: the Sallen & Key filter which transfer function is:

$$H(p) = 1/\left(\tau^2 p^2 + 2\zeta\tau p + 1\right)$$

The components (R1, R2, C) are related to the tuning parameters as:

$$f_c = \frac{1}{2\pi C\sqrt{R1R2}} \qquad \zeta = \frac{R1+R2}{2\sqrt{R1R2}}$$

These relations are straightforward but both tuning coefficients are coupled. It is therefore difficult to vary the one while the other remains constant. This structure is therefore not recommended when the parameters should be tuned dynamically and when low damping-factors are desired.

2.2.2 Digital design, Canonical

The canonical second order structure can be implemented as:

$$\begin{aligned} y(n) = {} & a_0 x(n) + a_1 x(n-1) + a_2 x(n-2) \\ & - b_1 y(n-1) - b_2 y(n-2) \end{aligned}$$

It can be used for any transfer function of the following type:

$$H(z) = \frac{a_0 + a_1 z^{-1} + a_2 z^{-2}}{1 + b_1 z^{-1} + b_2 z^{-2}}$$

In order to modify the cutoff frequency or the damping factor, all 5 coefficients have to be modified. They can be computed from the specification in the frequency plane or from a prototype analog filter. One of the methods that can be used is based on the bilinear transform [1]. The following set of formulas compute the coefficients for a lowpass filter:

f_c analog cutoff frequency
ζ damping factor
f_s sampling rate
$C = 1/[\tan(\pi f_c / f_s)]$

$$\begin{cases} a_0 = 1/\left(1 + 2\zeta C + C^2\right) \\ a_1 = 2a_0 \\ a_2 = a_0 \\ b_1 = 2a_0\left(1 - C^2\right) \\ b_2 = a_0\left(1 - 2\zeta C + C^2\right) \end{cases}$$

This structure has the advantage that it requires very few elementary operations to process the signal itself. It has unfortunately some severe drawbacks. Modifying the filter tuning (f_C, ζ) involves rather complex computations. If the parameters are varied continuously, the complexity of the filter is more dependant on the coefficient computation than on the filtering process itself. Another drawback is the poor signal to noise ratio for low frequency signals.

Other filter structures are available that cope with these problems. We will again review a solution in the analog domain and its counterpart in the digital domain.

2.2.3 State Variable Filter, analog

For musical applications of filters one wishes to have an independent control over the cutoff frequency and on the damping factor. A technique originating from the analog computing technology can solve our problem. It is called the state variable filter. This structure is more expensive than the Sallen & Key but has independent tuning components (R_f, R_ζ) for the cutoff frequency and the damping factors:

$$f_c = 1/\left(2\pi R_f C\right) \qquad \zeta = R/\left[2\left(R + R_\zeta\right)\right]$$

Furthermore it provides simultaneously three types of outputs: lowpass, highpass and bandpass.

2.2.4 State Variable Filter, digital

The state variable filter has a digital implementation as follows (Figure 2) [2].

$x(n)$ input signal
$y_l(n)$ lowpass output
$y_b(n)$ bandpass output
$y_h(n)$ highpass output

$$\begin{cases} y_l(n) = F_1 y_b(n) + y_l(n-1) \\ y_b(n) = F_1 y_h(n) + y_b(n-1) \\ y_h(n) = x(n) - y_l(n-1) - Q_1 y_b(n-1) \end{cases}$$

With tuning coefficients F1 and Q1, related to the tuning parameters f_C and ζ as:

$$F1 = 2\sin(\pi f_c / f_s) \quad Q1 = 2\zeta$$

it can be shown that the lowpass transfer function is:

$$r = F1 \quad q = 1 - F1Q1$$
$$H(z) = r^2 / \left[1 + \left(r^2 - q - 1\right)z^{-1} + qz^{-2}\right]$$

This structure is particularly effective not only as far as the filtering process is concerned but above all because of the simple relations between control parameters and tuning coefficients.

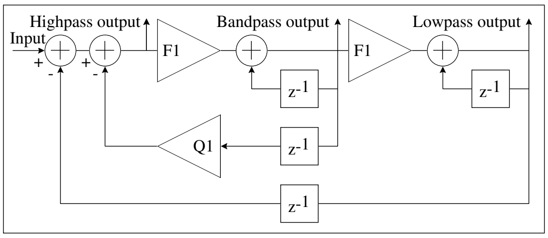

Figure 2: Digital state variable filter.

One should care for the stability of this filter, because at higher cutoff frequencies and larger damping factors it gets unstable. In most musical applications however it is not a problem because the tuning frequencies are usually small compared to the sampling frequency and the damping factor is usually set to small values [3, 4]. This filter has proven its suitability for a large number of applications. The nice properties for this filter have been exploited to produce endless glissandi out of natural sounds and to allow smooth transitions between extreme settings [5, 6].

We have considered here two different digital filter structures. More are available and each has its advantages and drawbacks. An optimum choice can only be made in agreement with the application [7].

2.2.5 Finite Impulse Response (FIR) Filters

The digital filter that we have seen before is said to have an Infinite Impulse Response (IIR). Because of the feedback loops within the structure, an input sample will excite an output signal whose duration is dependant on the tuning parameters and that can extent over a fairly long period of time. There are other filter structures without feedback loops. These are called Finite Impulse Response filters (FIR) because the response of the filter to an excitation sample lasts only a fixed period of time, however the filter is tuned.

These filters allow to build sophisticated filter types where strong attenuation of unwanted frequencies or decomposition of the signal in many frequency bands is necessary. They typically require more computing power than IIR structures to achieve similar results but when they are implemented in the form known as fast convolution they become competitive thanks to the FFT algorithm. It is rather unwieldy to tune these filters interactively.

As an example, let us briefly consider the vocoder application. If the frequency bands are fixed, then the FIR implementation can be most effective but if the frequency bands have to be subtly tuned by a performer, then the IIR structures will certainly prove superior [8].

2.3 Normalization

Filters are usually designed in the frequency-domain and we have seen that they have an action also in the time domain. Another correlated impact lies in the loudness of the filtered sounds. The filter might produce the right effect but the result might be useless because the sound has become too weak or too strong. The method to compensate for these amplitude variations is called normalization. Usual normalization methods are called L_1, L_2 and L_∞. L_1 is used when the filter should never be overloaded under any circumstances. This is most of the times overkill. L_2 is used to normalize the loudness of the signal. It is accurate for broadband signals and fits many practical musical applications. L_∞ normalizes actually the frequency response. It is best when the signal to filter is sinusoidal or periodical.

With a suitable normalization scheme the filter can prove to be very easy to handle whereas with the wrong normalization, the filter might be rejected by musicians because they cannot operate it.

The normalization of the state variable filter has been studied in [9] where several implementation schemes are proposed that lead to an effective implementation. In practice, a first order lowpass filter that processes the input signal will perform the normalization in f_c and an amplitude correction in $\sqrt{\zeta}$ will normalize in ζ (Figure 3). This normalization scheme allows to operate the filter with damping factors down to 10^{-4} where the filter gain reaches about 74 dB at f_c.

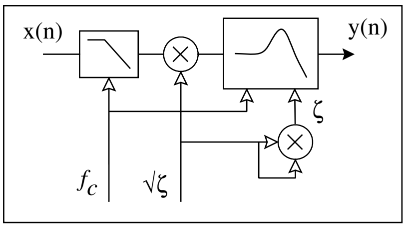

Figure 3: L_2-normalization in f_c and ζ for the state variable filter.

2.4 Sharp filters

Apart from FIR filters, we have given so far only examples of second order filters. These filters are not suitable for all applications. On the one hand, smooth spectral modifications are better realized by using first order filter. On the other hand, processing differently two signal-components that are close in frequency, or imitating the selectivity of our hearing system calls for higher order filters. FIRs can offer the right selectivity but again, they won't be easily tuned. Butterworth filters have attractive features in this case.

Such filters are optimized for a flat frequency response until f_C and yield a $-6n$ dB/octave attenuation for frequencies higher than f_C. Filters of order $2n$ can be built out of n second order sections, All sections are tuned to the same cutoff frequency f_C but each section has a different damping factor ζ (Table 1) [10].

n	ζ of second order sections				
2	0.707				
4	0.924	0.383			
6	0.966	0.707	0.259		
8	0.981	0.831	0.556	0.195	
10	0.988	0.891	0.707	0.454	0.156

Table 1: Damping factors for Butterworth filters.

These filters can be implemented accurately in the canonical second-order digital filter structure but modifying the tuning frequency in real time can lead to temporary instabilities. The state variable structure is less accurate for high tuning frequencies (i.e. $f_C > f_s/10$) but allows faster tuning modifications. A bandpass filter comprising a 4th order highpass and a 4th order lowpass was implemented and used to imitate a fast varying mute on a trombone [9]. Higher order filters (up to about 16) are useful to segregate spectral bands or even individual partials within complex sounds.

3 Delays

3.1 Acoustical delays

Delays can be experienced in acoustical spaces. A sound wave reflected by a wall will be superimposed to the sound wave at the source. If the wall is far away, such as a cliff, we will hear an echo. If the wall is close to us, we will notice the reflections through a modification of the sound color.

Repeated reflections can appear between parallel boundaries. In a room, such reflections will be called "flatter echos". The distance between the boundaries determines the delay that is imposed to each reflected sound wave. In a cylinder, successive reflections will develop at both ends. If the cylinder is long, we will hear an iterative pattern whereas, if the cylinder is short we will hear a pitched tone.

Equivalence to these acoustical phenomena have been implemented as signal processing units.

3.2 The FIR comb filter

The network that simulates a single delay is called the FIR comb filter. The input signal is delayed by a given time duration. The effect will be audible only when the processed signal is combined (added) to the input signal, which acts here as a reference. This effect has 2 tuning parameters: the amount of time delay τ and the relative amplitude of the delayed signal to that of the reference signal.

$$y(n) = x(n) + gx(n-M)$$
$$with\, M = \tau/f_s$$
$$H(z) = 1 + gz^{-M}$$

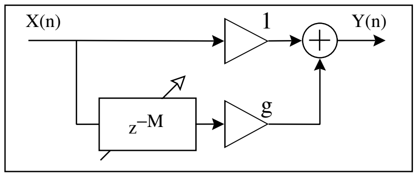

Figure 4: FIR comb filter.

The time response of this filter is made up of the direct signal and of the delayed version. This simple time domain behavior comes along with interesting frequency domain patterns. For positive values for g, the filter amplifies all frequencies that are multiple of $1/\tau$ and attenuates all frequencies that lie in between. The transfer function of such a filter shows a series of spikes and it looks like a comb. That is why this type of filter is called a comb filter. For negative values of g, the filter attenuates frequencies that are multiple of $1/\tau$ and amplifies those that lie in between. The gain varies between $1 + g$ and $1 - g$ [11].

As well as acoustical delays, the FIR comb filter has an effect both in the time domain and in the frequency domain. Our ear is more sensitive to the one aspect or to the other according to the range where the time delay is set. For larger values of τ, we can hear an echo that is distinct from the direct signal. The frequencies that are amplified by the comb are so close to each other than we barely identify the spectral effect. For smaller values of τ, our ear can no longer segregate the time events but can notice the spectral effect of the comb.

3.3 The IIR comb filter

Similar to the endless reflections at both ends of a cylinder, the IIR comb filter produces an endless series of responses $y(n)$ to an input $x(n)$. The input signal circulates in a delay line that is fed back to the input. Each time the signal goes through the delay line it is attenuated by g. It is sometimes necessary to scale the input signal by c in order to compensate for the high amplification produced by the structure. It is implemented as (Figure 5):

$$y(n) = cx(n) + gy(n-M)$$
$$with\, M = \tau/f_s$$
$$H(z) = c/\left(1 - gz^{-M}\right)$$

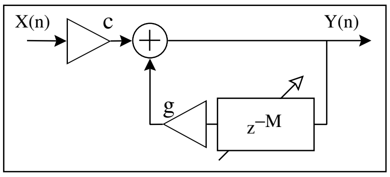

Figure 5: IIR comb filter.

Due to the feedback loop, the time response of the filter is infinite. After each time delay τ a copy of the input signal will come out with an amplitude g^p where p is the number of cycles that the signal has gone through the delay line. It can readily be seen, that $|g| <= 1$ is a stability condition. Otherwise the signal would grow endlessly.

The frequencies that are affected by the IIR comb filter are similar to those affected by the FIR comb filter. The gain varies between $1/(1-g)$ and $1/(1+g)$. The main differences between IIR comb and FIR comb is that the gain grows very high and that the frequency peaks get narrower as $|g|$ comes closer to 1.

3.4 Vibrato

When a car is passing by, we hear a pitch deviation due to the doppler effect [9]. This effect will be explained in another chapter but we can keep in mind that the pitch variation is due to the fact that the distance between the source and our ears is being varied. Varying the distance is, for our application, equivalent to varying the time delay. If we keep on varying periodically the time delay we will produce a periodical pitch variation. This is precisely a vibrato effect. For that purpose we need a delay line and a low-frequency oscillator to drive the delay time parameter. We should listen only to the delayed signal. Typical values of the parameters are: 5 to 10 ms as average delay-time and 5 to 14 Hz rate for the low-frequency oscillator [12].

3.5 Flanging, Chorus, Slapback, Echo

A few popular effects can be realized using the comb filter. They have become special names because of the peculiar sound effects that they produce. Consider the FIR comb filter. If the delay is in the range 10 to 25 ms, we will hear a quick repetition named "slapback" or "doubling". If the delay is greater than 50 ms we will hear an "echo". If the time delay is short (less than 15 ms) and if this delay time is continuously varied with a low frequency such as 1 Hz, we will hear the "flanging" effect. If several copies of the input signal are delayed in the range 10 to 25 ms with small and random variations in the delay times, we will hear the "chorus" effect (Table 2, [11, 12, 13], which is a combination of the vibrato effect with the direct signal. These effects can also be implemented as IIR comb filters. The feedback will then enhance the effect and produce repeated slapbacks or echos.

Delay range (ms) (Typ.)	Modulation (Typ.)	Effect name
0 ... 20	-	Resonator
0 ... 15	Sinusoidal	Flanging
10 ... 25	Random	Chorus
25 ... 50	-	Slapback
> 50	-	Echo

Table 2. Typical delay-based effects.

3.6 Phasing

The previous effects rely on the notches of the comb filter which are periodical in frequency. Another effect uses notches that are non periodical: the "phasing". A set of notch filters, that can be realized as a cascade of second order IIR sections, is used to process the input signal. The output of the notch filters is then combined with the direct sound. The frequencies of the notches are slowly varied using a low-frequency oscillator [14]. "The strong phase shifts that exist around the notch frequencies combine with the phases of the direct signal and cause phase cancellations or enhancements that sweep up and down the frequency axis" [11]. Although this effect does not rely on a delay line, it is often considered to go along with delay-line based effects because the sound-effect reminds that of flanging.

3.7 Normalization

We have seen in §2.3 that it is important to compensate for the intrinsic gain of the filter structure. Whereas in practice the FIR comb filter does not amplify the signal by more than 6 dB, the IIR comb filter can yield a very large amplification when $|g|$ comes close to 1. The L_2 and L_∞ norm are:

$$L_2 = 1\Big/\sqrt{1-g^2} \qquad L_\infty = 1/(1-|g|)$$

The normalization coefficient $c = 1/L_\infty$, when applied, ensures that no overload will occur with, for example, periodical input signals. $c = 1/L_2$ ensures that the loudness will remain approximatively constant for broadband signals.

3.8 Natural sounding resonators

We have made the comparison between acoustical cylinders and IIR comb filters. This comparison might seem inappropriate because the comb filters sounds "metallic". They tend to amplify very much the high frequency components and they appear to resonate much too long compared to the acoustical cylinder. To find an explanation, let us consider the boundaries of the cylinder. They reflect the acoustic waves with an amplitude that decreases with frequency. If the comb filter should sound like an acoustical cylinder, then it should also have a frequency dependant feedback coefficient *g(f)*.

This frequency dependance can be realized by using a first order lowpass filter in the feedback loop (Figure 6) [15].

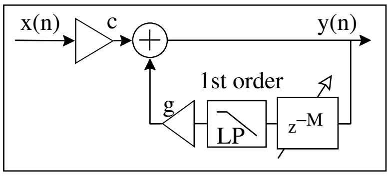

Figure 6: Lowpass IIR comb filter.

These filters sound more "natural" than the plain IIR comb filters. They find application in room simulators. Further refinements such as fractional delays and compensation of the frequency-dependant group-delay within the lowpass filter make them suitable for the imitation of acoustical resonators. They have been used for example to impose a definite pitch onto broadband signals such as sea waves or to detune fixed-pitched instruments such as a celtic harp [16].

4 Modulations

In the realm of telecommunications the word "modulate" means: "to shift the frequency spectrum of a signal to another frequency band". Numerous techniques have been designed to achieve this goal and some have found applications for digital audio effects.

4.1 Ring modulation

In the "ring modulation" (RM) the signal is multiplied by a carrier frequency. In the analog domain it was pretty difficult to do properly but within a computer it is straightforward since it is a mere multiplication. The input signal is called modulator *m(t)* and the second operand is called carrier *c(t)*:

$$y(t) = m(t) \times c(t)$$

If *c(t)* is a sine wave of frequency f_c, the spectrum of the output *y(t)* is made up of two copies of the input spectrum: the lower side band (LSB) and the upper side band (USB). The LSB is reversed in frequency and both sidebands are centered around f_c. Depending on the width of the spectrum of *m(t)* and on the carrier frequency, the side bands can be partly mirrored around the origin of the frequency axis. If the carrier signal comprises several spectral components, the same effect happens with each component. Whereas the audible result of a ring modulation is fairly easy to comprehend for elementary signals, it gets very complicated with signals having numerous partials. The carrier itself is not audible in this kind of modulation. When carrier and modulator are sine waves of frequencies f_c and f_m, one hears the sum and the difference frequencies $f_c + f_m$ and $f_c - f_m$.

4.2 Amplitude modulation

The amplitude modulation (AM) was easier to realize with analog electronic means and has been therefore in use for a much longer time. It can be implemented as

$$y(t) = [1 + \delta m(t)]c(t)$$

where it is assumed that the peak amplitude of *m(t)* is 1. The δ coefficient determines the depth of modulation. The modulation effect is maximum with $\delta = 1$ and the effect is disengaged when $\delta = 0$.

Typical application is with an audio signal as carrier *c(t)* and a low-frequency oscillator (LFO) as modulator *m(t)*. The amplitude of the audio signal varies according to the amplitude of the LFO, and it is heard as such. When the modulator is an audible signal and the carrier a sine wave of frequency f_c, the spectrum of the output *y(t)* is similar to that of the ring modulator except that the carrier frequency can be also heard. When carrier and modulator are sine waves of frequencies f_c and f_m, one hears three components: carrier, difference and sum frequencies ($f_c - f_m, f_c, f_c + f_m$).

One should notice that due to the integration time of our hearing system the effect is perceived in a different manner depending on the frequency range of the signals. A modulation with frequencies below 20 Hertz will be heard in the time domain (variation of the amplitude) whereas modulations by high frequencies will be heard as distinct spectral components (LSB, carrier, USB).

4.3 Single side band modulation

The upper and lower side bands carry the same information although organized differently. In order to save bandwidth and transmitter power, radio-communication engineers have designed the single side band (SSB) modulation scheme. Either the LSB or the USB is transmitted.

Phase shifted copies of the modulator and of the carrier are produced by Hilbert transformation. The upper and lower side bands can be computed as follows:

$m(t)$	modulator signal
$\hat{m}(t)$	Hilbert transform of $m(t)$
$c(t)$	carrier signal
$\hat{c}(t)$	Hilbert transform of $c(t)$
$USB(t)$	$= m(t)c(t) - \hat{m}(t)\hat{c}(t)$
$LSB(t)$	$= m(t)c(t) + \hat{m}(t)\hat{c}(t)$

This effect is typically used with a sine wave as carrier of frequency f_c. The use of a complex oscillator for *c(t)* simplifies the implementation. By using positive or negative frequencies it is then possible to select the USB or the LSB. The spectrum of *m(t)* is frequency shifted up or down according to f_c.

Out of harmonic sounds, this effect produces sounds that seem to be non-harmonic: a plucked-string sound is heard, after processing, like a drum sound. The modification in pitch is much less than expected, probably because our ear recovers pitch information out of the frequency difference between partials and not only out of the lowest partial of the tone [9].

5 Demodulations

Each modulation has a suitable demodulation scheme. The demodulator for the ring modulator uses exactly the same scheme, so no new effect is to expect there. The demodulator for the amplitude detector is called in the realm of digital audio effects an amplitude follower. Several schemes are available, some are inspired from analog designs, some are much easier to realize using digital techniques. These demodulators are made of three parts: a detector, an averager and a scaler.

5.1 Detectors

The detector can be a halfwave rectifier $d_h(t)$, a fullwave rectifier $d_f(t)$, a squarer $d_r(t)$ or an instantaneous envelope computer $d_i^2(t)$. The first two detectors are directly inspired from analog designs. They are still useful to achieve effects having typical analog behavior. The third and fourth types are much easier to realize in the digital domain. The four detectors are computed as follows:

$$\begin{aligned}
&x(n) && \text{input signal}\\
&\hat{x}(n) && \text{Hilbert transform of } x(n)\\
&d_h(n) && = \max[0, x(n)]\\
&d_f(n) && = \mathrm{abs}[x(n)]\\
&d_r(n) && = x^2(n)\\
&d_i^2(n) && = x^2(n) + \hat{x}^2(n)
\end{aligned}$$

5.2 Averagers

In the analog domain, the averager is realized with a resistor-capacitor (RC) network and in the digital domain using a first order lowpass filter. Both structures are characterized by a time-constant τ. The filter is implemented as:

$$\begin{aligned}
&g && = \exp[-1/(f_s\tau)]\\
&d(n) && \text{detector output}\\
&y(n) && = (1-g)d(n) + gy(n-1)
\end{aligned}$$

The time-constant must be chosen in accordance with the application. A short time-constant is suitable when fast variations of the input signal must be followed. A larger time-constant is better to measure the long-term amplitude of the input signal.

This averager is nevertheless not suited for many applications. It is often necessary to follow short attacks of the input signal. This calls for a very small time-constant, 5 ms typically. The output of the averager will then react very fast to any amplitude variation, even to the intrinsic variations within a period of a low frequency signal. We understand that we need an averager with two time-constants: an attack time-constant τ_a and a release time-constant τ_r. To distinguish it from the basic averager, we will name this one the "AR-averager". McNally has proposed an implementation having two fixed coefficients [17, 7] and Jean-Marc Jot an alternative where a single coefficient is varied according to the relation between the input and the output of the averager:

$$\begin{aligned}
&g_a = \exp[-1/(f_s\tau_a)] \qquad g_r = \exp[-1/(f_s\tau_r)]\\
&d(n) \quad \text{detector output}\\
&\begin{cases} \text{if } y_{ar}(n-1) < d(n) \text{ then } g = g_a \text{ else } g = g_r\\ y_{ar}(n) = (1-g)d(n) + gy_{ar}(n-1) \end{cases}
\end{aligned}$$

5.3 Scaling

The outputs of the systems described above are all different. In order to get measures that are comparable with each other, it would be necessary to scale the outputs. Although scaling schemes are typically defined for sine waves, each type of signal will require a different scaling scheme.

To build a RMS-detector or an instantaneous envelope detector, a root extractor would still be necessary, but building an accurate device can be difficult in analog and computationaly intensive in digital. Fortunately, it is often possible to avoid the root extraction by modifying the circuit that makes use of the averager output, so that it works fine with squared measures.

For these practical reasons the scaling is most of the times taken into account within the device that follows the averager output. If this device is a display, then the scaling can be done by changing the display marks.

5.4 Applications

Well known devices or typical applications relate to the previous schemes as follows:

- The AM-detector comprises the halfwave rectifier and of the basic averager.
- The Volume-Meter (VU-meter) is an AM-detector. It measures the average amplitude of the audio signal.
- The Peak Program Meter (PPM) is, according to DIN45406, a fullwave rectifier followed by an AR-averager with 10 ms integration-time and 1500 ms release-time.
- The RMS-detector, as found in electronic voltmeters, uses the squarer and the basic averager.

- A sound level meter uses a RMS-detector along with an AR-averager to measure impulsive signals.
- The RMS-detector associated with an AR-averager is the best choice for amplitude-follower applications in vocoders, computer music and live-electronics [18, 19, 20].
- Dynamics processors use various types of the above mentioned schemes in relation to the effect and to the quality that has to be achieved.
- The instantaneous envelope detector, without averager, is useful to follow the amplitude of a signal with the finest resolution. The output contains typically audio-band signals. A particular application of the $d_i^2(t)$ detector is the amplification of difference tones (Figure 7) [21, 22].

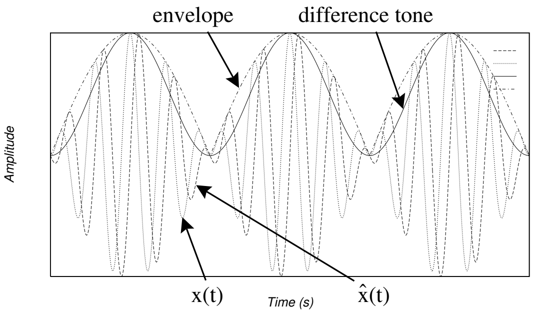

Figure 7: Instantaneous envelope detector as applied to detect a difference tone that is produced by two sine waves.

6 Conclusion

We have given here an overview but many details and applications had to be left aside. Only a few types of filters have been shown and many books could be filled with this topic [7, 4]. The issues that are particularly relevant to musical applications, such as normalization and smooth variation of the tuning parameters, have been stressed. The advantages of IIR filters have been emphasized but FIR filters have also fascinating applications; they could be reviewed along with time-frequency based audio-processing.

The delay-based algorithms are very popular but they are musically interesting only when enough attention is paid to the fine tuning of their parameters.

The envelope follower is most useful in its RMS version with independent attack and release times. The Hilbert transform has shown to be a decisive tool, as well for modulation as for demodulation.

The basic building blocks that have been reviewed here can be combined to build more complex digital audio effects.

7 References

[1] C. Dodge, T. A. Jerse: *Computer Music; Synthesis, Composition and Performance*. Schirmer Books, 1985.

[2] H. Chamberlin: *Musical Applications of Microprocessors*. Hayden Book Company, 1980.

[3] P. Dutilleux: *Simple to Operate Digital Time-Varying Filters*, Preprint 2757, 86th AES Convention, Hamburg, March, 1989.

[4] J. Dattoro: *Effect Design, Part 1: Reverberator and Other Filters*. J. Audio Eng. Soc., Vol. 45, No. 9, pp. 660-684, September, 1997.

[5] P. Dutilleux: *Spinning the Sounds in Real-Time*. ICMC Proc., Columbus, Ohio, November, pp. 94-97, 1989.

[6] P. Vasseur: *PURPLE FRAME*, in *Le sens caché 100 mystères*, CD FREE WAY MUSIQUE , n° 869193, 1993.

[7] U. Zölzer: *Digital Audio Signal Processing*. J. Wiley & Sons, 1997.

[8] M. Maiguashca: *Reading Castañeda.* 1 CD+ Booklet, Edition ZKM 3, Wergo2053-2, 1997.

[9] P. Dutilleux: *Vers la machine à sculpter le son, modification en temps réel des caractéristiques fréquentielles et temporelles des sons*. Thèse de l'Université d'Aix-Marseille II, 1991.

[10] M. Labarrère, J. Krief, B. Gimonet: *Le Filtrage et ses Applications*. CEPADUES, 1972.

[11] S. J. Orfanidis: *Introduction to signal processing*. Prentice Hall, 1996.

[12] C. Anderton: *Multieffects for Musicians*, Amsco Publications, 1995.

[13] J. Dattoro: *Effect Design, Part 2: Delay-Line Modulation and Chorus*. J. Audio Eng. Soc., Vol. 45, No. 10, pp. 764-788, October, 1997.

[14] J. O. Smith: *An Allpass Approach to Digital Phasing and Flanging*. ICMC Proc., pp 103-109, 1984.

[15] J. A. Moorer: *About This Reverberation Business*. Foundations of Computer Music, MIT Press, pp. 605-639, 1985.

[16] J.-C. Risset: *Luraï, pour Harpe celtique et bande*. Radio France, 1992.

[17] McNally G. W. *Dynamic Range Control of Digital Audio Signals*. J. Audio Eng. Soc., Vol. 32, No. 5 pp 316-327, May, 1984.

[18] P. Dutilleux: *Opéras multimédias, le rôle des ordinateurs dans trois créations du ZKM.* In *Musique et Arts Plastiques,* pp. 73-79. GRAME et Musée d'Art contemporain, Lyon, 1998.

[19] K. Furukawa: *"Swim, Swan"*. ZKM,1993.

[20] K. Furukawa: *Den ungeborenen Göttern*, Multimedia-Oper. ZKM, 1997.

[21] P. Dutilleux: *Verstärkung der Differenztönen (f2-f1)*. Bericht der 19. Tonmeistertagung Karlsruhe, pp. 798-806, Verlag K.G. Saur, 1996.

[22] M. Bach: *55 sounds for cello*. 1995.