Passband ripples and pre- and post-echoes of digital low-pass filters

Practical digital low-pass filters can have a perfectly linear phase characteristic, but they never have a perfectly flat magnitude response over their passband. Very often, they have ripples in their passband. Under some assumptions, these ripples can be described in terms of pre- and post-echoes.

The working assumption in this appendix is that there is some upper frequency f_{max} above which frequency components of sounds are inaudible, also when heard in combination with other frequency components, and that the cut-off frequency of the filter is greater than or equal to f_{max}. Hence, when an audio signal is passed through the filter, what happens above f_{max} has no impact on the sound of the output signal.

1. Simplification by assuming equal frequency distance between the ripples

Following R. Lagadec and T. G. Stockham [1], assume that the filter has linear phase and that its passband ripples have equal frequency distances. The filter can then be regarded as equivalent to a cascade of two linear-phase filters, a low-pass with perfectly flat passband and a filter that only has equidistant ripples all the way up to the Nyquist frequency. This second filter is the one that does the damage.

A filter with an impulse response consisting of a pre-echo, a main response and a post-echo produces exactly this kind of equidistant ripple response. That is, when the echoes occur a time T before and after the main signal and have a magnitude α, the transfer function (in the s domain) is

$$H(s) = 1 + \alpha e^{-sT} + \alpha e^{sT} \qquad \textbf{(B.1)}$$

neglecting the filter's constant delay for simplicity. Substituting $s = j\omega$, the corresponding magnitude response is

$$H(j\omega) = 1 + 2\alpha\cos(\omega T)$$

This produces ripples with a frequency distance of $1/T$.

It is interesting to see some numerical examples of the relation between the ripple magnitude and the magnitude of the pre- and post-echoes, see Table 1. Quite small ripples are required if you want to get the echoes down to a level that doesn't require temporal masking to make them inaudible. Fortunately, McClellan's FORTRAN program can readily produce filters with extremely small passband ripples.

Table 1 Some examples of ripple and echo magnitudes according to the theory of reference [36]

α	Echo magnitude	Ripple
0.03	-30.46 dB	+0.5061 dB / -0.5374 dB
0.005789	-44.75 dB	+0.09999 dB / -0.1012 dB
0.000403	-67.89 dB	+0.006998 dB / -0.007004 dB
0.000001	-120 dB	+0.00001737 dB / -0.00001737 dB

2. More realistic filter responses

The passband ripples of practical digital low-pass filters are not entirely evenly spaced in the frequency domain. Hence, it is unclear to what extent the theory of R. Lagadec and T. G. Stockham applies to realistic filters.

The pre- and post-echoes of a filter are, of course, determined entirely by its impulse response. Nonetheless, staring at the impulse response doesn't give us any idea of the echoes due to passband ripples, because the shape of the impulse response is mostly determined by what happens above the cut-off frequency.

The following method can make the impulse response aberrations due to passband ripples more visible:

1. Cascade the filter under test with an ideal continuous-time low-pass filter with cut-off frequency f_{max}. As we are interested in what happens below f_{max}, this doesn't affect the part of interest. When the passband is perfectly flat, the impulse response of the combined filter now has a $\sin(2\pi f_{max} t)/(2\pi f_{max} t)$ shape.
2. Sample the impulse response of the combined filter with a sample rate of 2 f_{max} and ensure that the samples fall on the zero crossings of the $\sin(2\pi f_{max} t)/(2\pi f_{max} t)$ function. Ideally, this produces only one nonzero sample, the one related to the main response. Any other peaks or $\sin(x)/x$ shaped ripples are due to the imperfect passband response of the filter under test.

Step 1 can be done mathematically by calculating the convolution of the impulse responses of the filter under test and the ideal filter. Step 2 means that you only need to calculate this convolution at a countable number of points. Overall, the calculation to be done is:

$$\sum_{k=0}^{N-1} \frac{\sin(2\pi f_{\max}(t-kT))}{2\pi f_{\max}(t-kT)} h(k) \qquad \textbf{(B.2)}$$

for

$$t=\frac{n}{2f_{\max}}+\left(\frac{N-1}{2}+\tau\right)T \qquad \textbf{(B.3)}$$

In these equations, N is the length and T is the sample period of the filter under test, $h(k)$ is its impulse response, f_{max} is the highest frequency of interest (as explained at the start of this appendix), t represents time and k and n are integers. τ is the number of sample periods that the main peak of the filter's impulse response is displaced from the centre; $\tau = 0$ for symmetrical filters.

As a typical example, I took a 128-tap, two-times-interpolating Parks-McClellan filter with 96 kHz output sample rate and a transition band from 21.768 kHz to 26.232 kHz. The weight was 1 in the passband and 12778 in the stopband, which resulted in a passband ripple of +0.007786 dB / -0.007793 dB and a stopband suppression of 143.08 dB (same order of magnitude as the filters used in the SRC4392). The echoes should theoretically have a level of -66.97 dB.

Using equations (B.2) and (B.3) with f_{max} = 21768 Hz, the plot of Figure 1 results. The largest echoes in the plot have a level of -69.22 dB. The difference with the theory of R. Lagadec and T. G. Stockham is less than 2.3 dB.

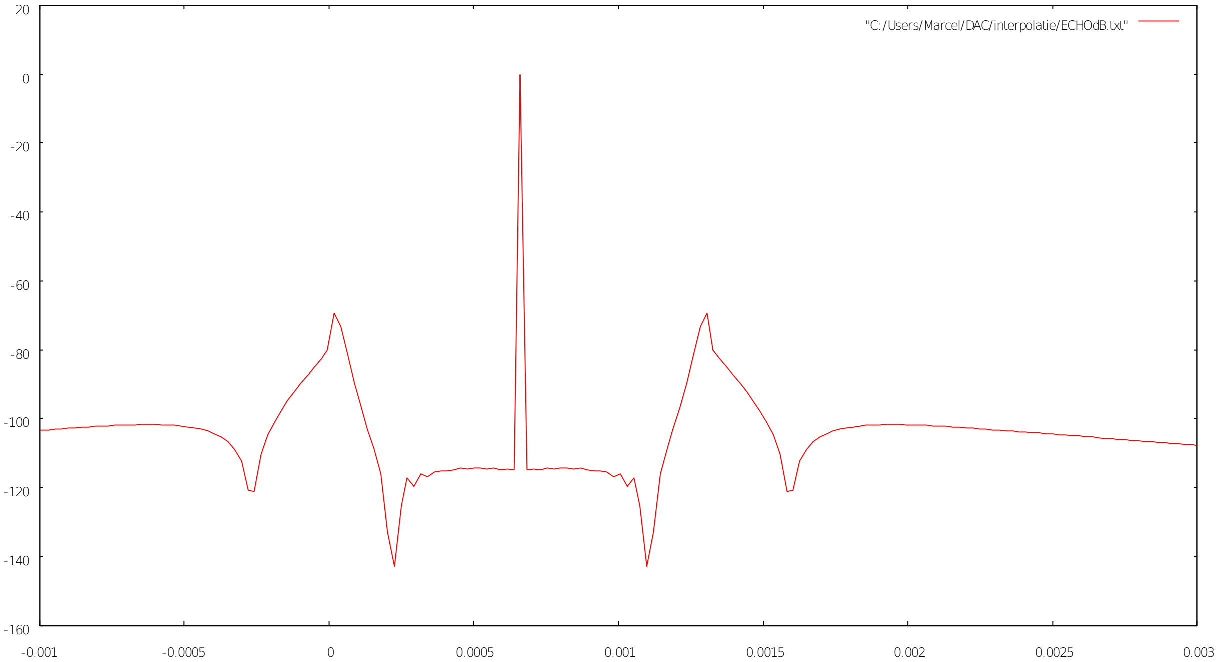

Figure 1 Impulse response aberrations due to the +0.007786 dB / -0.007793 dB passband ripple of a typical 128-tap filter. Horizontal scale: time points of equation (B.3) in seconds, vertical scale: result of equation (B.2) expressed in dB.

[1] R. Lagadec and T. G. Stockham, "Dispersive models for A-to-D and D-to-A conversion systems", Audio Engineering Society preprint 2097, presented at the 75th convention, March 1984