

Real-Time Filtering in BioExplorer

Filtering theory

Every signal can be thought of as built up from a large number of sine and cosine waves with different frequencies. A digital filter is a digital signal processor that attenuates a signal for each of its frequencies differentially. Ideally, the passband frequencies of the signal will be passed unchanged from input to output. The stopband frequencies, however, will be completely attenuated at the output of the filter. This implies an infinitely sharp fall-off of gain at the cut-off frequencies and a flat amplitude characteristic in both the passband and the stopband.

Unfortunately, such ideal filters do not exist. Two categories of filters are used to approximate the ideal behaviour, each with advantages and disadvantages: the Finite Impulse Response (FIR) filters and the Infinite Impulse Response (IIR) filters. The impulse response of a filter is the waveform that appears at the output of a filter when a Dirac-pulse (infinitely short unity pulse) is presented at the input.

FIR (Finite Impulse Response)

The FIR filter is a non-recursive filter: a filter that uses the current input sample of the signal together with previous input samples to calculate the current output sample.

$$y[n] = \sum_{k=0}^M b_k * x[n-k]$$

In this differential equation b_k are the impulse response coefficients. Because the output has to be calculated from a limited number of input samples, the impulse response is multiplied by a window of finite length (e.g. Hamming, von Hann and Kaiser). This yields the finite number of coefficients b_k (the order of the filter) that is necessary to calculate the output. However, applying this window has several disadvantages. The finite number of coefficients will not yield ideal filter characteristics. The fall-off will not be infinitely sharp, but some frequency bands around the cut-off frequencies will be partly attenuated, that is: some frequencies in the passband will exhibit a gain smaller than 1 and some frequencies in the stopband will exhibit a gain larger than 0. The fall-off will increase with increasing number of coefficients or, equivalent, order or window length. Furthermore, the behaviour in pass- and stopbands is influenced by the type of window used. The amplitude characteristic will not be flat in these bands, but will show a ripple. In the passband, the gain will not be 1 for every frequency, but each frequency will exhibit a slightly different gain. The larger the ripple, the larger the gain differences will be and the more the signal will degrade at the output. In the stopband, most frequencies that are supposed to be completely attenuated, will in some amount still be present in the output signal.

The length of the window will also determine the delay. When a signal is fed into a FIR filter, the output will have shifted its phase compared to the input. The phase shift is equal to half the length of the window, multiplied by the sample time; and each frequency in the signal will be shifted with the same constant phase shift.

FIR filter advantages

- A FIR filter will always give a stable output, because it is non-recursive
- A FIR filter has a delay that is constant over all frequencies; the signal shape will not be influenced by phase shifts

FIR filter disadvantages

- A FIR filter is inefficient, it uses lots of input to calculate the output, therefore:
- A FIR filter has a large delay

Practical implications

The use of a FIR filter is appropriate when the phase characteristic should be flat. When the demands on the filter are not very high, this is the easy solution to go for. In the context of EEG measurements this filter could be used as high pass or low pass filters. The EEG signal is not degraded by a variable phase characteristic and yet a sharp fall-off can be attained by picking a filter with a high order, since a large delay is allowed for the off-line application. These filters are especially well suited if the phase characteristic of the signal is very important and can thus be used for coherence and synchrony measures.

IIR (Infinite Impulse Response)

An IIR filter uses a different approach in calculating output. The IIR filter is recursive: it uses the current input sample together with previous output samples to calculate the current output sample.

$$y[n] = (N \sum_{k=1}^N) a_k * y[n-k] + (M \sum_{k=0}^M) b_k * x[n-k], N \geq 0, M > 0$$

The order of the filter is determined by the number of previous output samples used to calculate the current output sample. As was the case with FIR filters, increasing the order of the filter will improve the amplitude characteristics of the filter.

Naturally, there is a drawback to increasing filter order. More samples are used to calculate output, therefore the filter is less efficient. This results in a degradation of the phase characteristic.

The mechanisms underlying IIR filtering are beyond the scope of this discussion.

IIR filter advantages

- An IIR filter is very efficient, it uses only a few previous output values to calculate the current output, therefore:
- An IIR filter has only a small delay

IIR filter disadvantages

- An IIR filter can become unstable, if coefficients are not well chosen
- Because the requirement of causality and an infinite impulse response that converges to 0 for $n \Rightarrow \text{infinity}$, the impulse response cannot be symmetrical. This implies a phase characteristic that is not constant over frequencies. Therefore, different frequencies will appear at the output with a different shift in phase. The shape of the output signal will change as a consequence of different phase shifts.

Practical implications

The use of an IIR filter is appropriate when, a small absolute delay is required or when strong requirements on amplitude characteristic are necessary.

Where there is only one type of FIR filter, there are several types of IIR filters. Three of these will be discussed below: the Butterworth, the Chebyshev and the Elliptic IIR filters.

Butterworth IIR

A Butterworth filter is a filter which has a maximally flat characteristic in the passband: the gain decreases gradually from the center of the passband towards the -3 dB point. No ripple is present in the stopband either.

Practical implications

The use of a Butterworth IIR filter is appropriate when a maximally flat pass- and stopbands are important, when less stringent restrictions on fall-off and phase characteristics are allowed. Choice to put emphasis on fall-off or phase characteristics can be tuned by choosing an appropriate filter order. The phase characteristic of a Butterworth filter is poor, therefore synchrony or other phase related analysis are not recommended with this filter.

An example of the use of Butterworth filters in EEG would be an on-line EEG frequency band analysis (theta, gamma, alpha, etc). The pass- and stopbands are maximally flat, thus resulting in a quality output signal for the different frequency bands. Furthermore, a reasonably sharp fall-off can be reached. Care has to be taken in selecting the correct trade-off between fall-off and phase characteristic.

Chebyshev IIR

A Chebyshev filter has a constant ripple in the passband which oscillates between 1 and $(1+\epsilon^2)^{-1/2}$; ϵ can be tuned. The number of ripples in the passband increases with filter order. There is a trade-off between the behaviour in passband and stopband. In comparison with a Butterworth filter, the stopband behaviour and fall-off of a Chebyshev are better.

Practical implications

The use of a Chebyshev IIR filter is appropriate when some ripple in the passband is allowed and more emphasis is placed on good stopband behaviour and fall-off. The phase characteristic of a Chebyshev filter is poor. The use of this filter is recommended with intermediate frequency ranges e.g. alpha 1 vs. alpha 2.

Elliptic IIR

An elliptic IIR filter has a ripple of constant height in both passband and stopband, the number increasing with filter order. In comparison with the Butterworth as well as the Chebyshev, at a given order, the fall-off is better for an elliptic filter.

Practical implications

The use of an elliptic IIR filter is appropriate when the fall-off is the very critical design criterion.

The sensori-motor rhythm (SMR) is a good example of such a critical design criterion. The SMR is a very specific rhythm of $\sim 12-14$ Hz. For a quantification of the SMR –and only the SMR– a very narrow frequency band has to remain at the filter output, thus implying a very steep fall-off. Some waveform degradation has to be accepted (ripple, phase characteristic) to attain this sharp fall-off, but this is a small price to pay knowing that you analyze a true sensorimotor rhythm.

Filter setting possibilities in BioExplorer

Band range: the range of frequencies at which the -3dB points can be placed, -3dB point meaning a gain of the input signal of 0.5 at that particular frequency.

Order/length: the order of a filter is determined by the number of values (input in the case of FIR, input & output in the case of IIR) that are used to calculate the output sample.

Ripple: a ripple is an oscillation on the stop- and passbands in the amplitude spectrum, ideally these would be flat.

FIR

Band range: 0.0 – 60.0 Hz
Order/length: 5 – 400
Ripple: n.a.
Attenuation: n.a.

Butterworth IIR

Band range: 0.0 – 60.0 Hz
Order/length: 1 – 8
Ripple: n.a.
Attenuation: n.a.

Chebyshev IIR

Band range: 0.0 – 60.0 Hz
Order/length: 1 – 8
Ripple: min. 0.001
Attenuation: n.a.

Elliptic IIR

Band range: 0.0 – 60.0 Hz
Order/length: 1 – 8
Ripple: min. 0.001
Attenuation: -200 – -10

Evaluation of filter spectra in BioExplorer

- In the following evaluations of filter spectra, the passband was set to [0.0 – 30.0]
- In BioExplorer the smallest possible –3dB point can be set to 0.0 Hz. This means that DC recording will always suffer from signal distortion, because the true DC level will exhibit a gain of 0.5. Thus, effectively, only half of the true DC level would remain at the output of the filter. The cut-off determines how much of the low frequency signal is attenuated. A very sharp cut-off will only attenuate the true DC level. A less sharp cut-off will attenuate low frequencies as well (with a gain of 0.5 –1; the higher the frequency, the less attenuated).
- Theoretically, the attenuation of the true DC level should always be –3dB, when a low cut-off of 0.0 Hz is chosen. However, the gain of the DC level shows a dependence on filter order and describes a cycle of a gain of 0 to 1. The number of orders that comprise a cycle seems to be dependent on the passband that is chosen.

FIR

Magnitude

- The cut-off becomes sharper with increasing order.
- Theoretically, the attenuation of the true DC level should always be –3dB, when a low cut-off of 0.0 Hz is chosen. However, the gain of the DC level shows a dependence on filter order and describes a cycle of a gain of 0 to 1. The number of orders that comprise a cycle seems to be dependent on the passband that is chosen.
- There seems to be a small ripple on the passband, which –according to theory– is equal to the ripple in the stopband. This ripple is independent of filter order.

Delay

- The delay is constant over all frequencies and increases with increasing order.

The choice of filter order is a matter of finding the right balance between fall-off and delay.

Butterworth IIR

Magnitude

- The cut-off becomes sharper with increasing order.
- At higher orders the amplitude spectrum flattens out nicely.
- Filters with an even order display bad behaviour at the lower frequencies (gain = 0 or gain = 1.4).
- The passband ripple (and stopband ripple???) becomes smaller at higher filter orders.
- Order 7 seems to yield the best amplitude spectrum.

Delay

- The delay (at orders > 1) is much larger in the vicinity of the cut-off frequencies, as it is at the middle frequencies.
- Order 5 seems a good option here: small delay at DC level, somewhat bigger (0.035) in the 3-20 Hz plateau, and ascending to 0.06 in the 20-30 Hz range.

Amplitude characteristics seems OK at order 5; order 5 would be my choice for DC measurements with a Butterworth filter.

Chebyshev IIR

Magnitude

- When a minimal ripple is set, fall-off increases with increasing filter order. The behaviour in the middle frequencies is really good at higher orders. Low frequencies, however, exhibit very bad behaviour; either gain is too high or too low.
- If a larger ripple is allowed, the behaviour in the middle frequencies deteriorates. On the other hand there is an improvement in the behaviour at lower frequencies: the very high or very low gain is only present at the very low frequencies close to DC level. Effectively, this also means that the cut-off is improved when a larger ripple is allowed.

Delay

- When the ripple is set to minimal, the phase spectrum is very bad. At some filter orders, very high plateaus comprising some frequency intervals rise up to a large delay. Actually, the delay is acceptable at no order, except order 1 & 2.
- At larger ripples, delay properties improve a lot. At higher orders the delay is much higher at the cut-off frequencies, but the fall-off to the middle frequencies (small delay plateau) is reasonably sharp. The absolute delay, however, is fairly large at orders higher than order 4.

At a ripple of 1 (which actually seems too large) order 4 would be a good option. When a smaller ripple is required (e.g. 0.3), order 7 seems a good choice.

Elliptic IIR

Magnitude

- At a minimal ripple and attenuation the behaviour at the middle frequencies is very acceptable and of course again the sharper cut-off at higher filter orders. The difference in gain, however, is very big between low and middle frequencies.
- If a larger ripple is allowed, the same result as with the chebyshev is attained: sharper fall-off at low and high cut-off frequencies.
- With a minimal ripple and maximal attenuation (-10, 0) the amplitude spectrum exhibits acceptable but worse behaviour at the middle frequencies (at higher filter orders). The fall-off is very sharp, but displays oscillations in the vicinity of the cut-off frequency. This leads to unacceptable behaviour at the lower frequencies. Furthermore, the stopband shows large oscillations.
- If a larger ripple is allowed, the behaviour in the passband is very good, especially at orders 3, 5 and 7 (ripple = 1, att. = -10). However, the stopband shows large ripples.

Delay

- (min. ripple, min. att.) The delay shows very bad phase spectra with big jumps in the spectra. Order 8 seems OK, but with a very large delay at the DC level. Order 5 and 6 are better, but without the flat spectra.

- (ripple = 1.0, min. att.) Phase spectra improve when ripple is allowed larger. At higher orders, there is a plateau of low delay at the middle frequencies, with a fairly sharp fall-off at low and high cut-off frequencies. The 4th order has an acceptable phase spectrum, which is flat for lower and middle frequencies (reasonable amplitude spectrum too).
- (min. ripple, max. att.) Higher order filter have large jumps in the phase spectrum in different frequency bands (larger delay in bands closer to DC when filter order is higher, bands become smaller with increasing order).
- (ripple = 1.0, att. = -10) At higher orders very good phase spectra are obtained, with a very steep fall-off at low and high cut-off frequencies. The absolute value of the delay is hard to estimate, due to the large scale, but seems to be fairly small.

An advisable setting would be: Order 8, ripple = 0.3, att. = -40