Design and Optimization of Anti-Aliasing Filters for Ocean Turbulence Measurements

August 2005
Rolf Lueck

Introduction

Aliasing is an endemic phenomenon when sampling data at equally spaced intervals. Signal components above the Nyquist frequency will appear at lower frequencies in the sampled signal, and this will result in an ill-formed signal spectrum. To ensure unbiased measurements it is necessary to suppress, as much as possible, the signal energy above the Nyquist frequency *before* the signal is sampled (or digitized).

This article examines the effects of aliasing using a synthetic signal as an example, and discusses how data acquisition systems can be optimized to suppress aliasing effects. The final section of the article discusses the practical application of the anti-aliasing filters, as they are used in instruments for the measurement of oceanic turbulence, manufactured by RGL Consulting.

An example of aliasing

Except for the quantum decorrelation time of photons, signals generated by natural processes have unlimited bandwidth. It is possible to have any or all frequency components represented in the signal. However, sampling a signal at regular intervals at a rate f_S samples per second, the signal's frequency content is limited to the Nyquist frequency, which is

$$f_N = \frac{1}{2} f_S \ . \tag{1}$$

Unfortunately, any frequency components in the original (continuous) signal that are above the Nyquist frequency will appear at a frequency smaller than the Nyquist frequency in the sampled signal. This effect is called aliasing, which can lead to a misrepresentation or distortion of the sampled signal's spectrum.

The aliasing effect is easily demonstrated by sampling a synthetic sinusoidal waveform, such as

$$\sin\left(2\pi kt\right) = \sin\left(\frac{2\pi kn}{128}\right), \ n = 0,1,...127 \tag{2}$$

where the left-hand side of (2) represents the continuous domain signal (defined for all time, *t*) and the right-hand side represents the sampled signal, at 128 samples per second. The signal frequency is *k* [Hz] and the Nyquist frequency is f_N = 64 Hz. Figure 1 shows a representation of the sampled signal for selected frequencies, *k*. The sampled values are shown as dots. The straight line between the samples is a visual aid only and represents nothing because we have no information about the continuous signal (in the sampled domain).

When the frequency, *k*, is small compared to the Nyquist frequency (Figure 1, blue line), the sampled signal looks like a sine wave and the plot of the sampled signal is a good representation of the continuous domain.

When the signal frequency increases to one-half of the Nyquist frequency (Figure 1, green line) the samples no longer look like a sine wave. The observed frequency is still correct but we are sampling the signal at its peaks, troughs and zero-crossings. The straight line segments mislead the eye by making the waveform look like saw teeth.

When we increase the frequency of the signal to the Nyquist frequency (Figure 1, red line), every sample is zero because samples are taken at the zero-crossing of the signal. If we had, instead, sampled a cosine wave, then the samples would have alternated between the peaks and troughs of the wave with a frequency exactly equal the frequency, *k*, of the continuous signal.

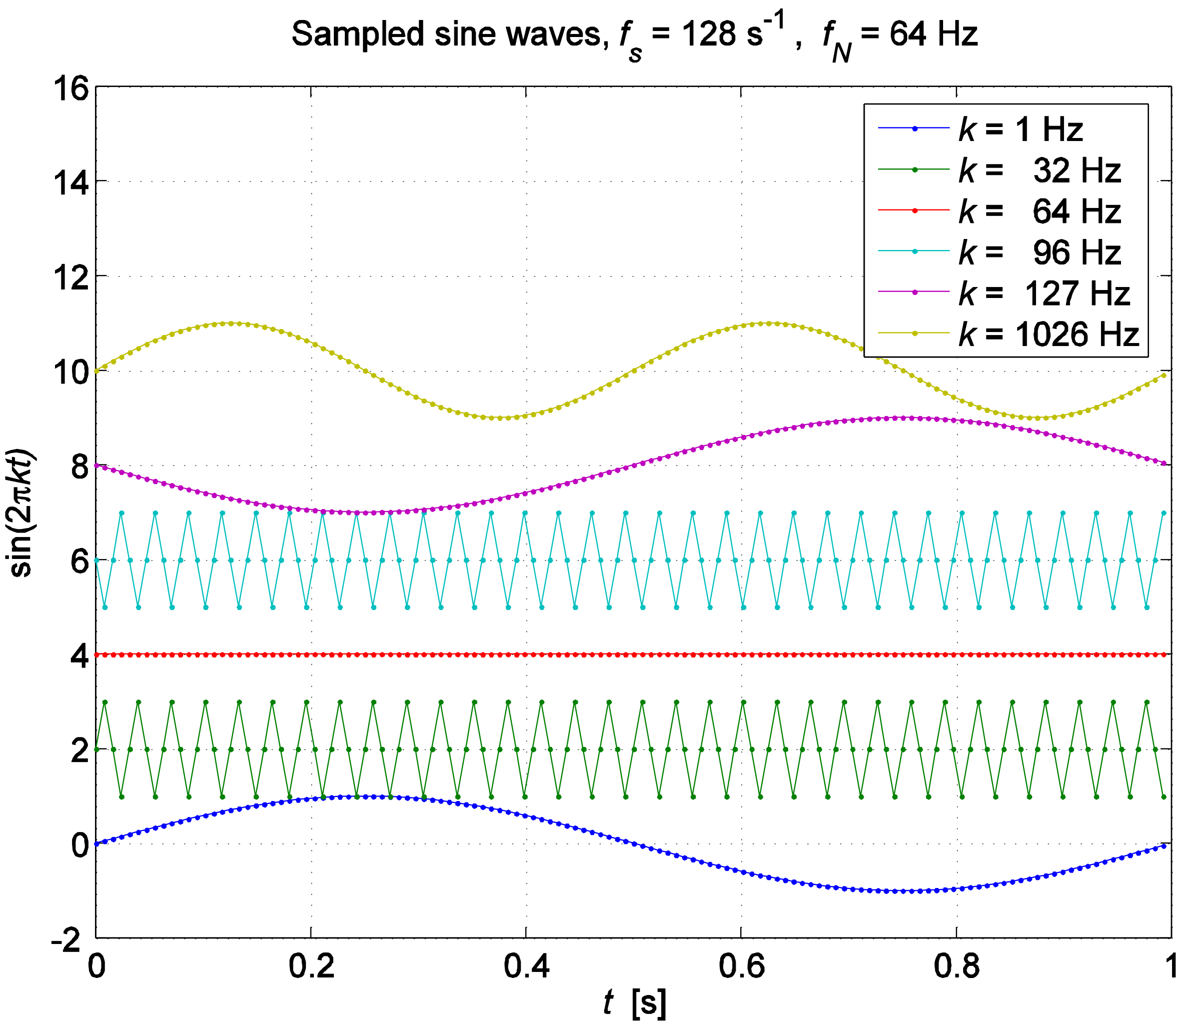

Figure 1. Sinusoidal waves of frequency *k* sampled at 128 s^{-1}. Each frequency case is progressively offset by 2 for visual clarity.

A further increase of the frequency to 1.5 f_N (k = 96 Hz, Figure 1, cyan line) gives a peculiar result. The waveform is a phase inverted version of the sampled signal when k = 32 Hz. The apparent frequency of the sampled signal is wrong. It has been aliased downward to a frequency $\widehat{k} = 2 f_N - k$. The same shifting of frequency and phase inversion is seen when k = 127 Hz, which results in a sampled waveform with the proper frequency (1 Hz) but with inverted phase (Figure 1, magenta line). Had we sampled a cosine wave, then we would have had the same frequency shifting but without the phase inversion.

Finally, there is no frequency limit to aliasing. A continuous domain signal with the ridiculously high frequency of 1026 Hz looks like a waveform of 2 Hz, after it is sampled (Figure 1, yellow line).

In summary, any signal component with frequency above the Nyquist frequency will be detected in the sampling process. An “alias” of that signal component will appear at a frequency below the Nyquist frequency. The aliased signal may also appear phase shifted.

Designing anti-aliasing filters

To avoid aliasing effects during the data acquisition, it is necessary to include an electronic low-pass filter in the data acquisition system. The low-pass filter must be designed to eliminate the frequency content above the Nyquist frequency *before* the digitization of the signal. This is conceptually simple, but the practical application is not straight forward. There are no low-pass filters that either completely stop all signals above its pass-band or have an infinitely steep transition from the pass-band to stop-band.

The most commonly used filter in data acquisition systems is the Butterworth type because it is easily implemented electronically. Its amplitude response is (Figure 2, green line)

$$H_B(f) = \frac{1}{\sqrt{1+\left(\frac{f}{f_C}\right)^{2n}}}\,, \tag{3}$$

where f_C is the cut-off frequency and n is the filter order. The higher the order, the sharper the transition from pass-band to stop-band. The practical limit for the filter order is about $n = 8$. The attenuation of a Butterworth filter is monotonic with increasing frequency (Figure 2 and Figure 3).

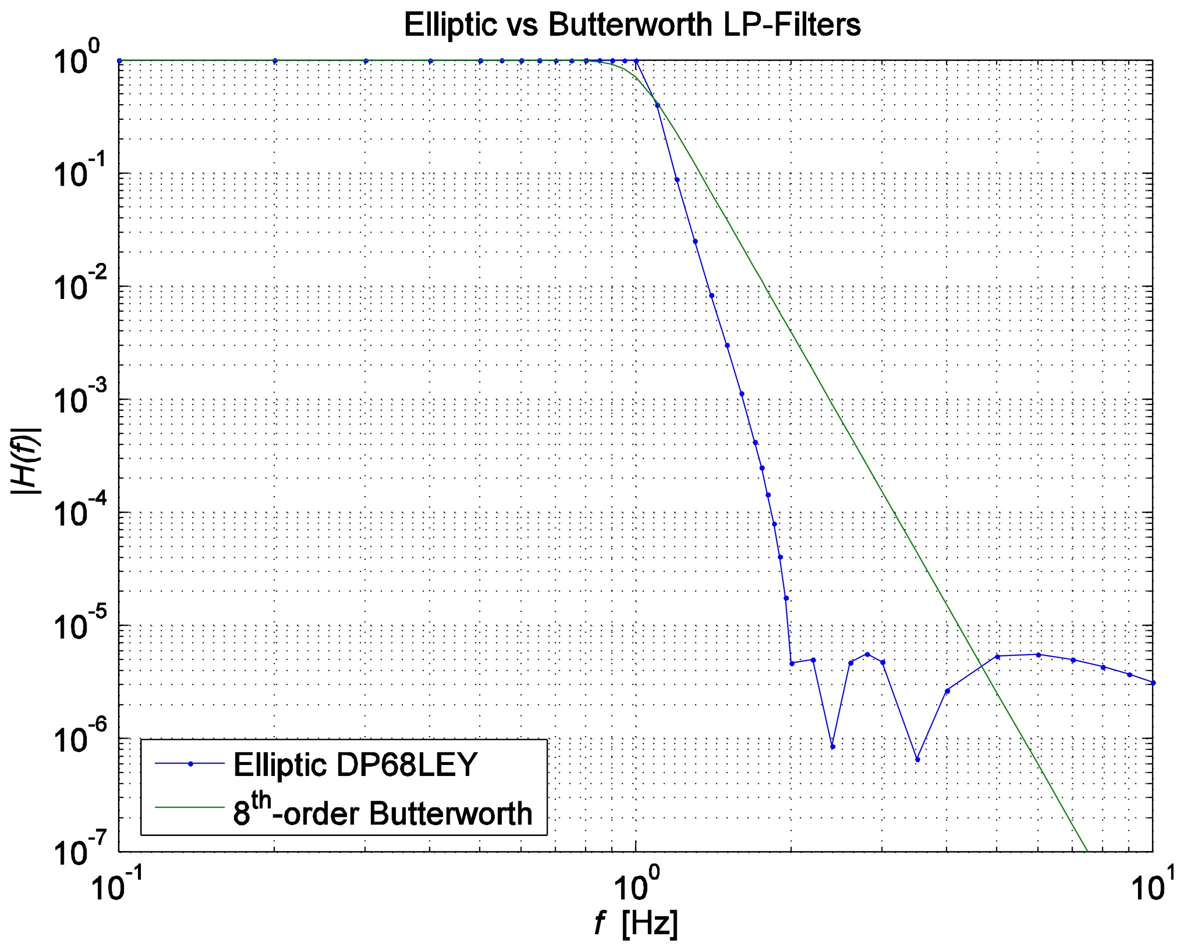

Figure 2. The amplitude frequency response of an 8th –order Butterworth filter commonly found on commercial data acquisition systems (green line) and that of Elliptic filter modules used with the RGL Consulting filter boards (blue line) and. The frequency is normalized to a cut-off frequency of $f_c = 1\ \mathrm{Hz}$.

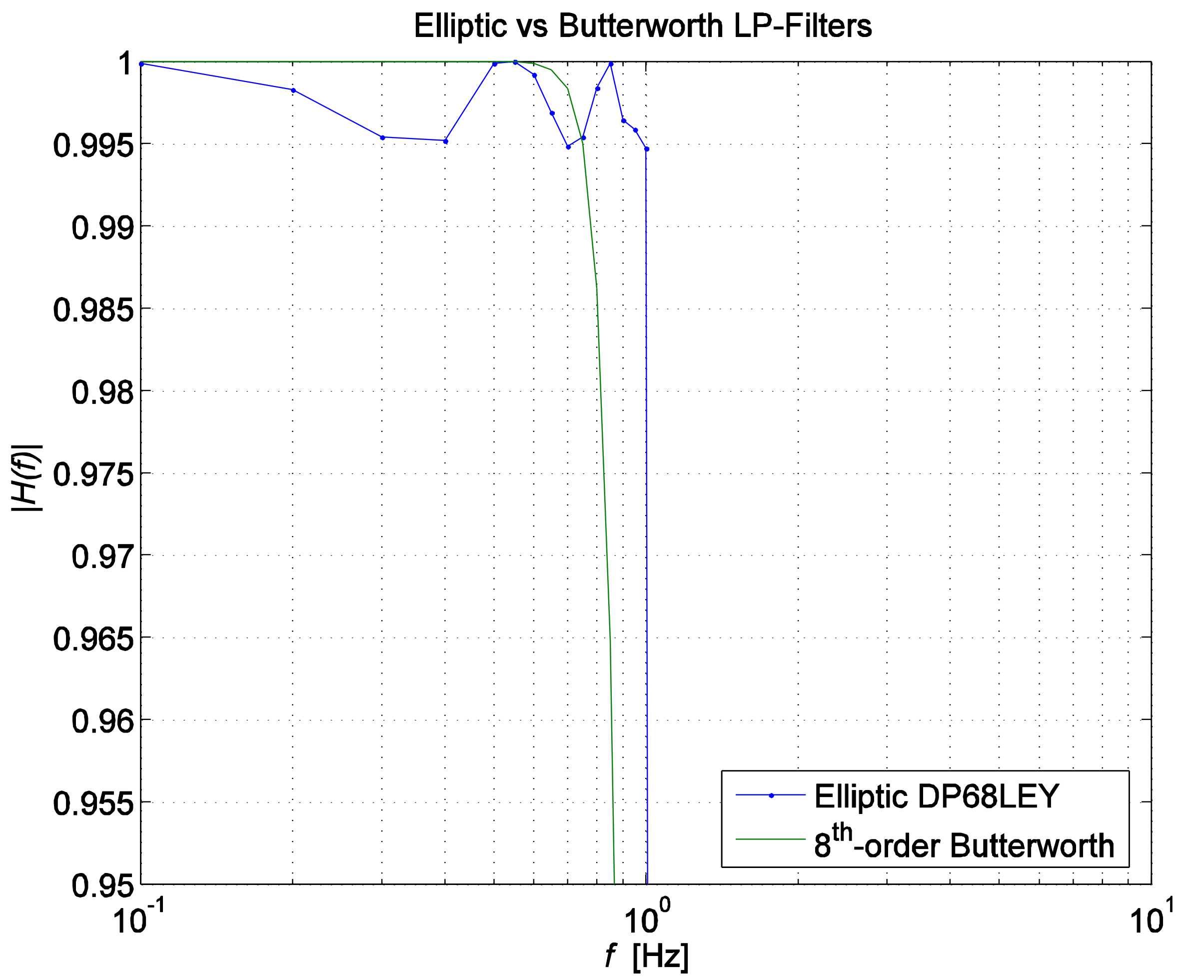

Figure 3. Same as Figure 2 but with a close up view of the pass-band.

The question of how much attenuation is required to avoid aliasing is, unfortunately, not easily answered. We will assume that we have a 16-bit data sampling system (such as the one found on instruments by RGL Consulting). The theoretical limit of the standard deviation of the sampling noise of an ideal analog-to-digital converter is $1/\sqrt{12} = 0.29$ when expressed in terms of quantization step size (or bit-resolution). The maximum amplitude of a sampled signal is $\frac{1}{2}2^{16} \approx 30000$, so one estimate of the attenuation required to suppress aliasing is the ratio of these two numbers, namely, 10^5. (However, signals with amplitude smaller than 0.29 may still be detected if they have very narrow bandwidth.) The Butterworth filter reaches this level of attenuation for frequencies more than 4 times larger than its cut-off frequency (Figure 2).

The pass-band of the filter must be equal to, or larger, than the “band-of-interest”, which is defined by the scientific interest of the user of the data acquisition system. For this example, we will set this band-of-interest to 1 Hz so that all other frequencies are normalized with respect to this band. Using a Butterworth filter, we must then set the Nyquist frequency to be 4 Hz in order to achieve the factor 10^5 in signal reduction. Consequently, the sampling rate must be 8 Hz. We see that the finite transition rate of the Butterworth filter requires a rather high sampling rate, leading to larger data storage requirements (for recording systems) and/or higher telemetry rates (for real-time data transmission).

One alternative to Butterworth filters are elliptic filters (Figure 2, blue line). The use of elliptic filters helps to reduce the sampling rate relative to the band-of-interest. Elliptic filters have the sharpest transition of all filters but this comes at a price. The pass-band is not flat but has small “ripples” (Figure 3), and the asymptotic attenuation in the stop-band is finite (Figure 2). Elliptic filter modules are also more expensive than Butterworth filters, as the design and construction of elliptic filters is beyond the capability of most electronic manufacturers.

RGL Consulting uses modules produced by Frequency Devices Inc. Transition band sharpness can be traded off against high-frequency attenuation and pass-band ripple. Filters sharper than the DP68LEY do exist but they have less attenuation in the stop band and are, therefore, not suitable for 16-bit data acquisition systems.

In our example the elliptic filter reaches an attenuation of 10^5 at 2 Hz and so a sampling rate of 4 samples per second is sufficient to avoid aliasing. This is a factor of two improvement over the Butterworth filter.

In summary, a sampled signal can be rendered free of aliased signals by filtering the continuous domain signal with an appropriate low-pass filter before sampling it. For a given frequency band-of-interest, the sharpness of the filter's transition band determines the required sampling frequency. Conversely, for a given sampling rate, the filter's transition determines the maximum useful frequency in the sampled signal.

Optimizing anti-alias filters

So far, we have assumed that all frequencies from 0 to the Nyquist frequency must be free of aliasing. This is a very conservative approach. We can relax this requirement without compromising the signal fidelity if we ensure that only the band-of-interest be free of aliasing rather than all frequencies smaller than the Nyquist frequency.

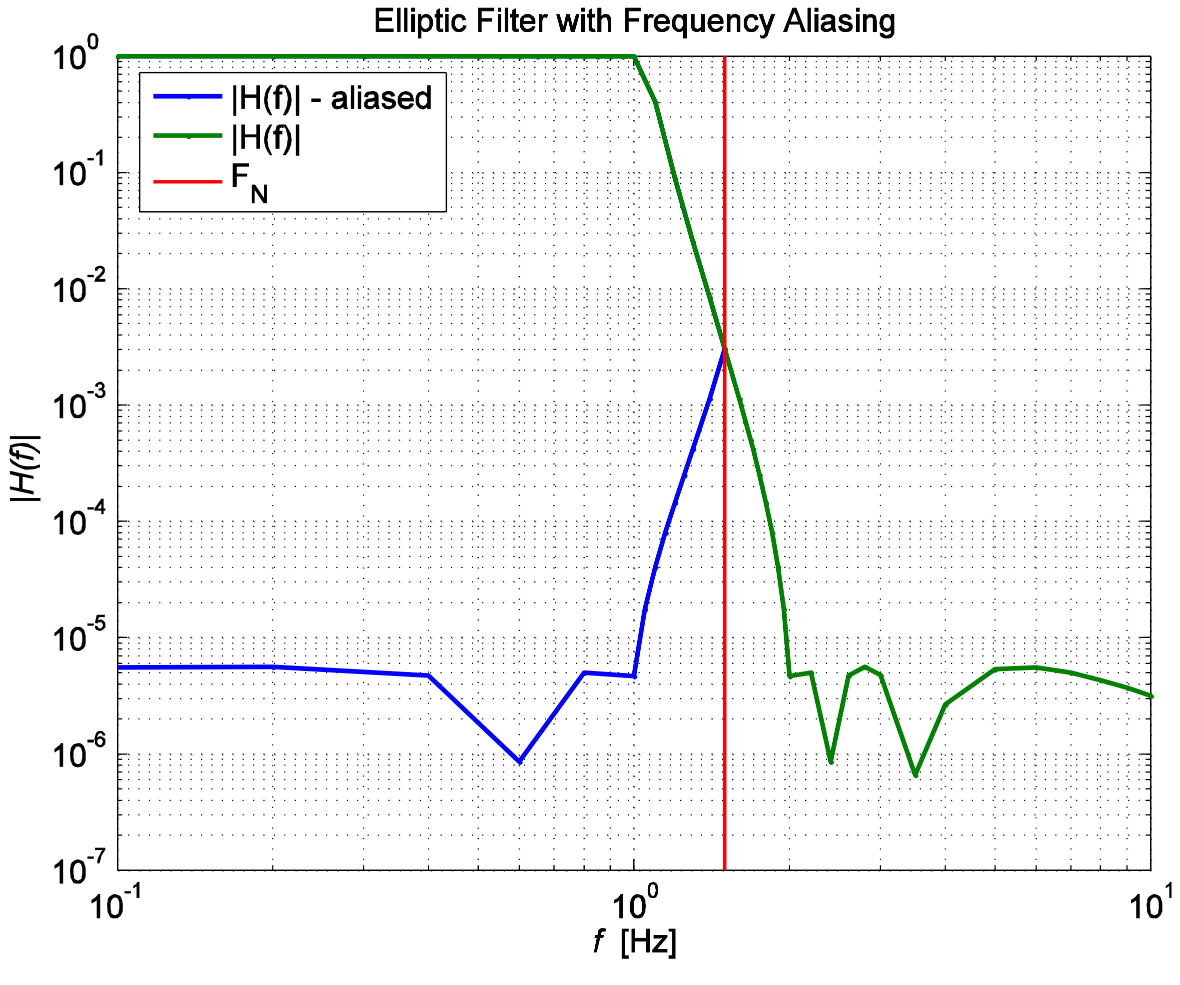

Figure 4. The attenuation of the Elliptic filter (green curve) and the attenuation of continuous domain signals which are aliased back into the Nyquist band (blue curve).

The elliptic filter has a 2-to-1 transition range. If we place the Nyquist frequency exactly mid-way between f_C and $2f_C$. We have the mathematical relationship that

$$2f_C - f_N = f_N - f_C. \quad (4)$$

This means that

$$\frac{f_c}{f_N} = \frac{2}{3}. \tag{5}$$

Signals that have frequencies between f_N and $2f_C$ can be aliased back into the Nyquist band, but they will not reach the band-of-interest. The frequency components that are aliased back into the band-of-interest are attenuated by more than 10^5 (Figure 4, blue line). This means that the sampling rate can be as low as $3f_C$ with an Elliptic filter.

A practical application

The measurement of oceanic small-scale turbulence requires the rapid sampling of highly variable signals of fluid velocities or changes in scalar properties (e.g., temperature). The band-of-interest for oceanic turbulence measurements depends on the speed at which the sensors move through the water column. For a profiling speed, *U*, the band of interest is approximately $100U$ cycles per meter. Most instruments travel slower than 1 $\mathrm{m\,s^{-1}}$, so the band-of-interest is usually smaller than 100 Hz.

For the data analysis, sampling rates that are integer powers of 2 are convenient. The lowest power-of-2 sampling rate that is more than 3 times larger than 100 Hz is $512\ \mathrm{s^{-1}}$. Working backwards from this sampling rate dictates that $f_C < 170\ \mathrm{Hz}$. In our instruments we chose $f_C = 165\ \mathrm{Hz}$ to give the filter a margin of safety (Figure 5). If we reduce the sampling rate to only 256 samples per second, then the cut-off frequency must be reduced to 83 Hz. This may be too small for instrument profiling faster than $0.8\ \mathrm{m\,s^{-1}}$ and leaves no margin for future sensors that may have better spatial resolution than 100 cpm. Thus, we chose a sampling rate of 512 $\mathrm{s^{-1}}$ and a cut-off of $f_C = 165\ \mathrm{Hz}$ to give users wide latitude with respect to profiling speed and to accommodate future developments.

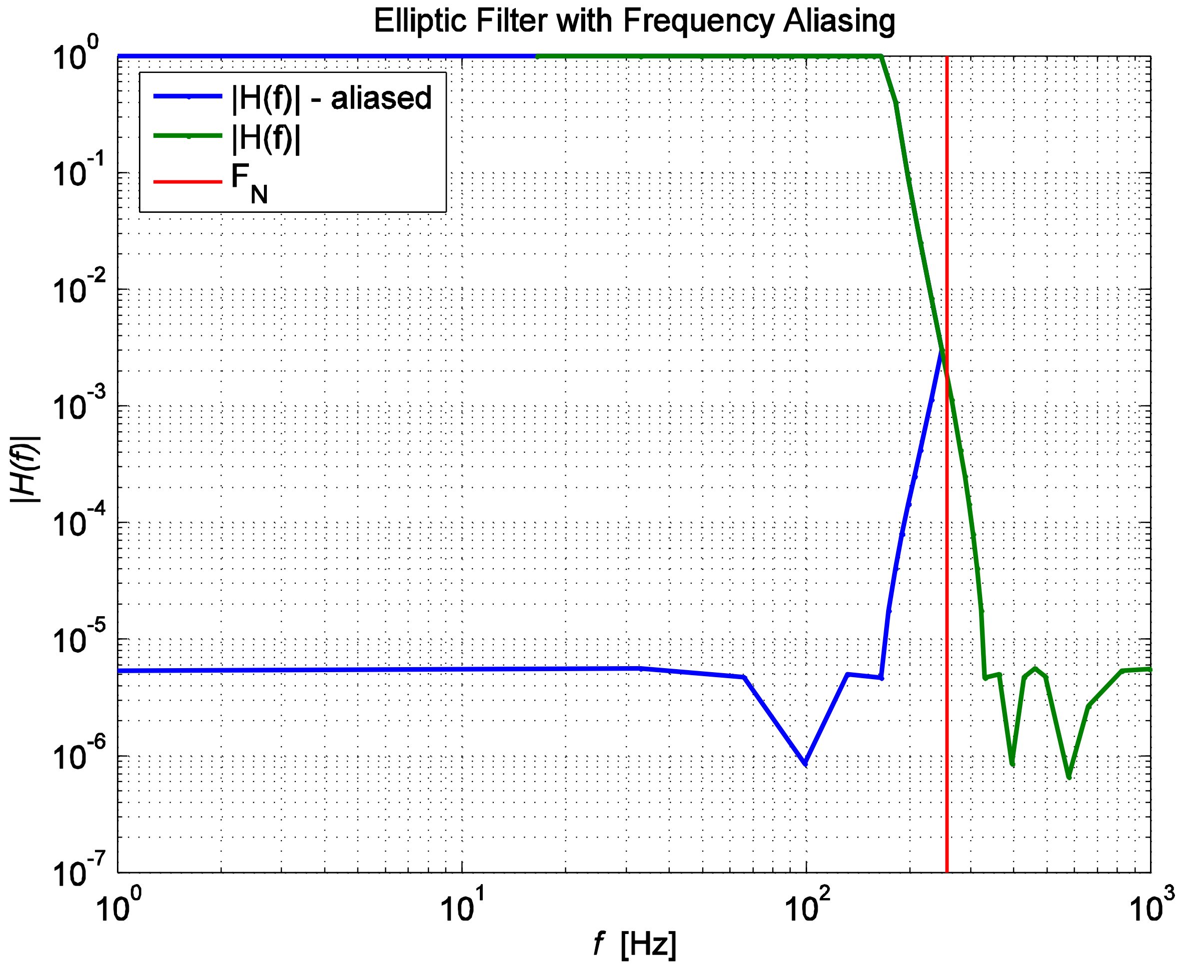

Figure 5. Same as Figure 4 but scaled to a band-of-interest of 165 Hz to show the actual filter response and potential aliasing with a typical instrument manufactured by RGL Consulting ltd.

End-of-Document