RSI Technical Note 010

# Design and Optimization of Anti-Aliasing Filters

Rolf Lueck

2018-01-03

Rockland Scientific International Inc.
520 Dupplin Rd
Victoria, BC, CANADA, V8Z 1C1
www.rocklandscientific.com

# Contents



# List of Figures

# 1 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, a signal's energy above the Nyquist frequency *before* it is sampled (or digitized).

This note illustrates the effects of aliasing using a synthetic signal (section 2). The limitations of sampling are discussed in section 3 because of their implications for an anti-aliasing filter. The requirements of an anti-aliasing filter are quantified in section 4 and the implicit bandwidth limitation of alias-free samples is presented in section 5. The RSI filter design is found in section 6. This note concludes with brief advice on determining if an anti-aliasing filter was used in a sampled time series.

# 2 An example of signal aliasing

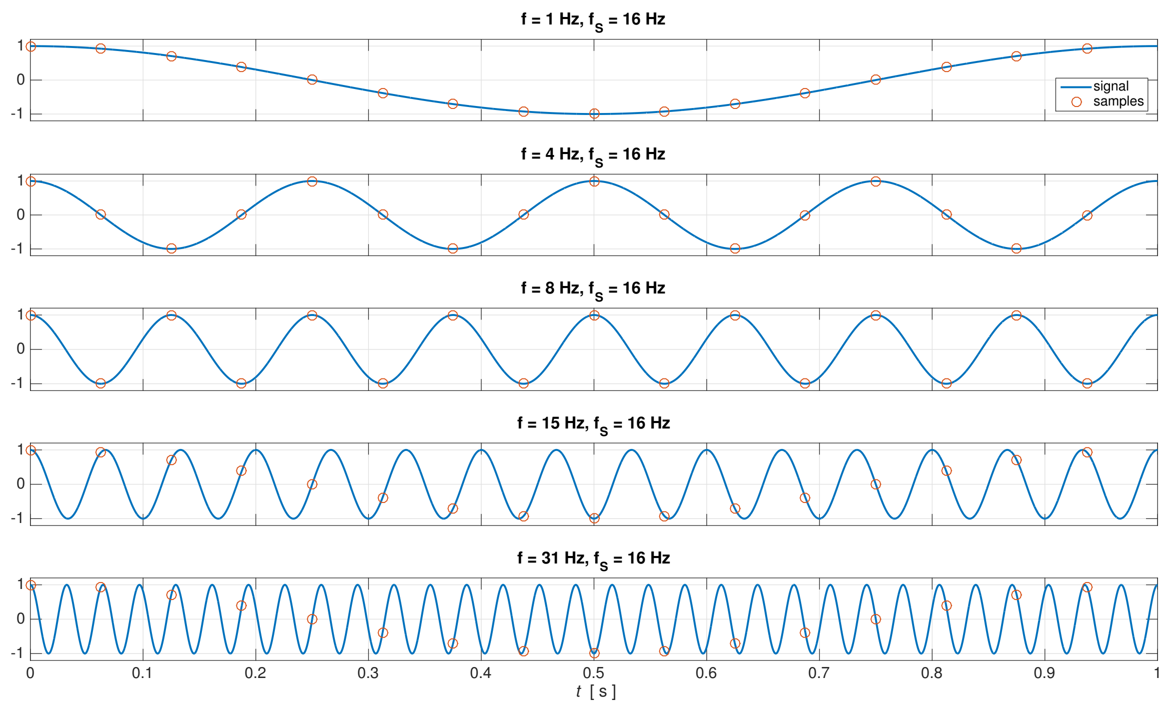


Figure 1: An example of signal aliasing. In each panel, the blue line is the continuous domain signal $\cos(2\pi f t)$ with the frequency indicated above each panel. The cosine signal is sampled (red circles) at a rate of $f_s = 16\,\mathrm{s}^{-1}$ ($\Delta = 0.0625\,\mathrm{s}$) in all panels.

One way to visualize aliasing is to sample a cosine oscillation at a fixed rate, $f_s = 1/\Delta$. For example, $f_s = 16\,\mathrm{s}^{-1}$ (Figure 1). Samples of a 1 Hz cosine oscillation are not ambiguous (upper panel). When the frequency is raised to 4 Hz (second panel from the top), we get only 4 samples per cycle, but the frequency of the signal is still not ambiguous. Even if the frequency is increased to the Nyquist frequency, $f_N = f_s/2$, (third panel), we sample the peaks and troughs of the waveform. However, if the phase had been shifted by $\pi/2$, we would have sampled the zero crossings, seen nothing, and possibly concluded that there was no signal. When the frequency is raised to 15 Hz (second panel from bottom), our sampling clearly shows a cosine oscillation of 1 Hz and, we get the identical samples when the frequency is 31 Hz (bottom panel). In fact, we would see a 1 Hz cosine oscillation for any frequency $nf_s - 1$, where $n$ is any whole number. Aliasing has no bounds.

The message is clear, we must sample a signals at a rate that is at least twice the highest frequency in the signal. This amounts to a 'chicken and egg' problem. How do we determine the highest frequency in a signal, so that we may sample it at an adequate rate, other than to actually sample it, which provides an ambiguous result because of aliasing? The practical solution to this problem is to construct the measurement system so that the signal is *intentionally* low-pass filtered to remove (or at least strongly attenuate ) the signal for all frequencies higher than the Nyquist frequency. This part of the continuous domain processing is usually the last stage of the signal conditioning. These filters are called *anti-aliasing* filters to reflect their purpose.

Before we can quantitatively specify or choose an anti-aliasing filter, we must first examine some limitations of sampling.

# 3 Data quantization

Up to here, we have been rather vague about the value of a sample and implicitly assumed that it is identical to the value of the signal at the moment of sampling. This is unrealistic. A physical signal (temperature, pressure, velocity, etc.) is continuous in both time *and* value. For example, if the temperature is decreasing from 4 °C to 3 °C, then it will at some moment during that cooling have a temperature of $\pi$ °C.

The sampled data must be represented by a finite number of digits or bits. This means that it is quantized into discrete levels. For example, if a system measures a signal, $s(t)$, over a range from $s_l$ up to $s_u$, and the values are quantified by a $B$-bit number, then there are $2^B$ quantum levels for the samples, and the step size of the quantum levels is

$$\delta_s = \frac{s_u - s_l}{2^B - 1} \approx \frac{s_u - s_l}{2^B} \; . \tag{1}$$

The difference, $s_u - s_l$, is usually called the full-scale range of a measurement system. The number of steps in a staircase is one less than the number of levels and, if $B > 10$, the difference between these two numbers is inconsequential.

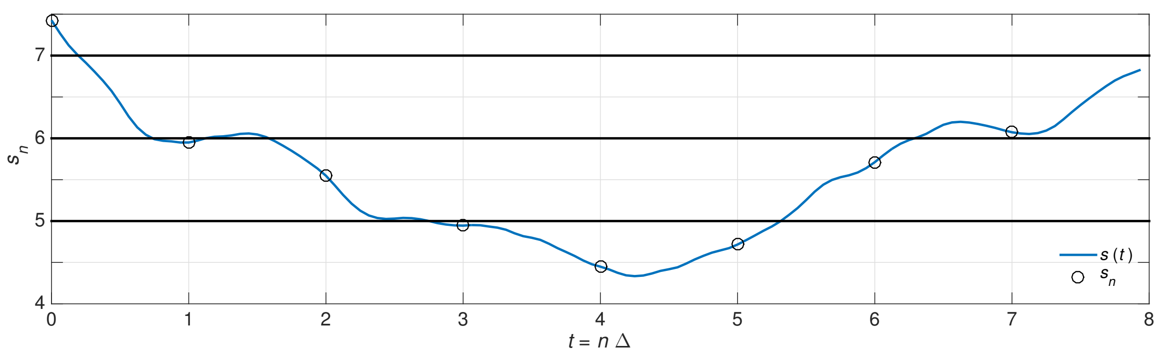


Figure 2: An example of the quantization of the samples, $s_n$, of a continuous signal $s(t)$. The ideal sampler assigns the nearest quantum level (black horizontal lines) to the samples.

The ideal sampler will assign the measured value to the nearest quantum level which would produce the data [7, 6, 6, 5, 4, 5, 6 and 6] for the example in Figure 2.

It is now quite common for the sampler (which is also called an analog-to-digital converter, or ADC) to provide 16-bit values. The step size (1) is then about $1.5 \times 10^{-5}$ of the full-scale range of the instrument. Some samplers provide up to 24-bit values.

Because the data are quantized, every sample is wrong but, ideally, by no more than $\pm\delta_s/2$. Good samplers get within a factor of 2 of the ideal quantization. The samples are, effectively, the true value of the signal plus random noise. Because the physical signal is continuous in value, the error of the samples taken with an ideal sampler must be uniformly distributed over the range of $\pm\delta_s/2$. The probability density function of such a uniformly distributed random variable is

$$\begin{aligned} p(x) &= \frac{1}{\delta_s} \quad |x| \le \delta_s/2 \\ &= 0 \quad\; |x| > \delta_s/2 \; . \end{aligned} \tag{2}$$

The mean error is the first moment of the probability density function

$$\mu = \int_{-\infty}^{\infty} \frac{1}{\delta_s} x \mathrm{d}x = \frac{1}{\delta_s} \int_{-\delta_s/2}^{\delta_s/2} x \ \mathrm{d}x = 0 \tag{3}$$

and so the ideal sampler does not bias the data. This means, among other things, that the average of your samples converges to the true mean of the signal. However, the variance of the error

$$\sigma^2 = \int_{-\infty}^{\infty} \frac{1}{\delta_s} x^2 \mathrm{d}x = \frac{1}{\delta_s} \int_{-\delta_s/2}^{\delta_s/2} x^2 \ \mathrm{d}x = \frac{2}{\delta_s} \frac{1}{3} \left( \frac{\delta_s}{2} \right)^3 = \frac{\delta_s^2}{12} \tag{4}$$

is not zero. Your data are the actual values of the signal *plus* uniformly distributed noise with a standard deviation of $\delta_s/\sqrt{12}$. This additional noise is called the *sampling noise*, or the *quantization noise*. The sampling noise cannot have any particular frequency and, therefore, it must be white and uniformly distributed over the Nyquist band, which means that the spectrum of this noise is $\phi_{\delta_s} = \delta_s^2/(12 f_N)$, where $f_N = f_s/2$ is the Nyquist frequency.

Your estimate of the spectrum of a signal is never smaller than the sampling noise. For example, the vibration spectra of Figure 3 flatten out to the sampling noise of $7 \times 10^{-4}\,\mathrm{counts}^2/\mathrm{Hz}$ [1] , near the Nyquist frequency, even though the continuous domain signals are attenuated to a lower level by the anti-aliasing filters. For the example of Figure 3, the ideal sampling noise (magenta) is

$$\phi_{\delta_s} = \frac{\delta_s^2}{12 f_N} = \frac{1}{12 \times 256} = 3.3 \times 10^{-4}\,\mathrm{counts}^2/\mathrm{Hz} \tag{5}$$

which is only two times smaller than the actual sampling noise (black).

The sampling noise variance is independent of the rate of sampling, and the noise spectrum decreases with increasing sampling rate. This is the main benefit of over sampling a signal.

[1] The data samples are just integer numbers and are, therefore, dimensionless. However, it is common to present values in units of counts to indicate explicitly their lack of dimensionality.

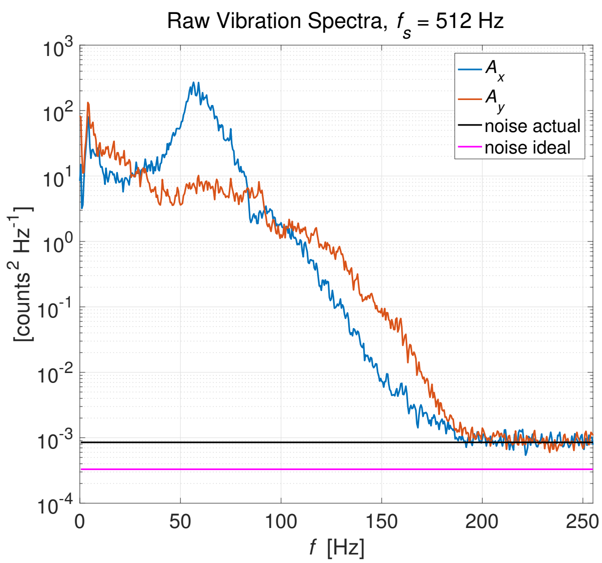


Figure 3: The spectra from two vibration sensors, $A_x$ and $A_y$, that have not been converted into physical units. The signals were sampled at $512\,\mathrm{s}^{-1}$. The anti-aliasing filter was a cascade of two 4-pole, low-pass, Butterworth filters with an effective cut-off frequency of 98 Hz.

# 4 The minimum requirements of an anti-aliasing filter

We can now quantify, in a semi-object manner the attenuation requirements of an anti-aliasing filter. The bandwidth of the analog-domain portion of a data sampler is astonishingly broad, typically up to several mega-hertz, which means that the sampler will see signals (and noise) up to very high frequency. This is true regardless of the rate of sampling! Therefore, it is important to strongly attenuate a signal at high frequency even if you are sampling at very modest rates. We will consider two cases – frequency independent noise, and an unwanted sinusoidal signal above the band of interest.

## 4.1 Frequency-independent noise

If the unwanted signal is frequency-independent continuous-domain noise, then we want the variance of this noise, after it is filtered, to be small compared to the sampling noise of $\delta_s^2/12$. That is, we want

$$\sigma_N^2 \ A^2 \ll \frac{\delta_s^2}{12} \tag{6}$$

where $\sigma_N^2$ is the noise variance and $A$ is the attenuation of the filter. If the noise is approximately Gaussian, then its standard deviation is approximately one-sixth of its peak-to-peak fluctuation. The peak-to-peak fluctuation is at most comparable to the full-scale range of the sampler, $2^B \delta_s$, and we can use this value for the maximum possible noise. Therefore,

$$A \ll \frac{\delta_s}{2\sqrt{3}} \frac{1}{\sigma_N} = \frac{\delta_s}{2\sqrt{3}} \ \frac{6}{2^B \ \delta_s} = 2^{-B}\sqrt{3} \ . \tag{7}$$

For a 16-bit sampler, with $B = 16$, $A$ should be small compared to $2.5 \times 10^{-5}$. The elliptic anti-aliasing filter of Figure 4 (blue) meets this requirement at a frequency that is 2 times higher than the cut-off frequency. However, the transfer function of this elliptic filter does not asymptotically go to zero. For this reason, a Butterworth filter (Figure 4, green) is often preferred for anti-aliasing, if a less steep attenuation with respect to frequency is acceptable.

## 4.2 Sinusoidal noise

If the unwanted signal is a sinusoidal oscillation, then the peak-to-peak amplitude of this oscillation, $2\alpha$, is at most comparable to the full-scale range of the sampler. Therefore, the maximum value of the periodogram, at the frequency of oscillation, is

$$\psi_S(f) = \frac{1}{2} A^2 \alpha^2 \ T \ . \tag{8}$$

where $T$ is the length, in seconds, of the data segments used in each fast Fourier transform (FFT) to construct the periodogram. We would like this contribution to the spectrum to

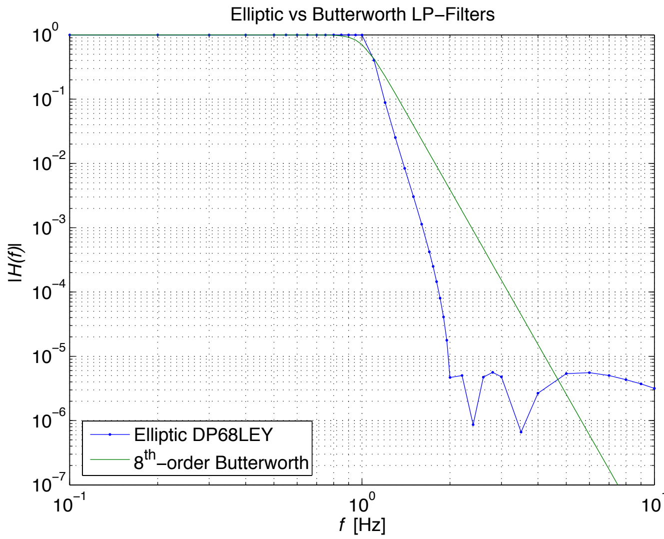


Figure 4: An example of an anti-aliasing filter with a cut-off frequency of 1 Hz.

be small compared to the spectrum of the sampling noise, $\psi_N(f) = \delta_s^2/(12 f_N)$. Therefore,

$$\begin{aligned} \psi_S(f) &\ll \psi_N(f) \\ \frac{1}{2}A^2\alpha^2\,T = \frac{1}{2}A^2\left(\frac{\delta_s\,2^B}{2}\right)^2 T &\ll \frac{\delta_s^2}{12 f_N}\ . \end{aligned} \tag{9}$$

This requires the attenuation, $A$, to be

$$A \ll \frac{1}{2^B}\sqrt{\frac{2}{3T f_N}} = \frac{1}{2^B}\sqrt{\frac{4}{3N}} \tag{10}$$

because $T f_N = N/2$, where $N$ is the number of samples used in each FFT. This dependence on $N$ makes it impossible to specify an attenuation, $A$, for all situations. For example, $\sqrt{N}$ is typically 10 to 100, i.e. not of order unity. For $B = 16$, the numerical factor in (9), excluding $N$, equals $1.7 \times 10^{-5}$. An FFT length of 1024 requires an attenuation small compared to $5 \times 10^{-7}$. An anti-aliasing filter that asymptotically goes to zero is also desirable for sinusoidal interference. The attenuation of such a filter increases with increasing frequency and only signals near the Nyquist frequency are problematic.

The two cases discussed in this section are “worst-case” examples, and the numerical values are quite conservative. For example, if the broad-band noise, or the amplitude of an unwanted sinusoidal signal, is close to the full-scale range of a measurement system, then there is no room left for your signal of interest, and you have a very poor instrument.

# 5 Implicit bandwidth limitation

Filters that are well suited for anti-aliasing purposes are those that have a small range of frequency over which their transfer function transitions from 1 to a very small value. The Butterworth family of filters are often used for anti-aliasing. The inverse of the magnitude-squared of their transfer function is

$$|H(f)|^{-2} = 1 + \left(\frac{f}{f_c}\right)^{2n} \tag{11}$$

where $f_c$ is the 'cut-off' frequency and $n$ is the order of the filter. The cut-off frequency is also known as the half-power frequency because the magnitude-squared response is one-half at that frequency. Even for moderate orders, the attenuation is quite strong for $f \gg f_c$ (Figure 4, green). There are many other types of filters suitable for anti-aliasing the data before sampling. For example, the sharpest transition from the pass- to the stop-band is provided by elliptic filters (Figures 4, blue).

Because the transition from passing to attenuating by a filter spans a finite range of frequency, it is impossible for a measurement system to resolve signals up to the Nyquist frequency, without aliasing. There is considerable confusion among people taking measurements regarding what is resolved in a signal. For example, it is often (and carelessly) stated that "if I sample at a rate of $f_s$, I will resolve frequencies up to $f_s/2$." This is an unrealistic expectation – the resolved range is significantly smaller than $f_s/2$.

For some instruments, the continuous-domain signal is averaged over the interval between samples. It is foolish to think that this averaging eliminates aliasing. Such averaging is identical to data smoothing by convolution with a box-car (sometimes called a uniform or top-hat) weighting function. This kind of smoothing does reduce the spectrum of a signal above the sampling rate but it is a very poor filter compared to Butterworth (and most other) filters. Signal aliasing by this sort of sampling (i.e. the averaging of samples) is invariably worse than that achieved by the proper filtering of a signal before sampling.

# 6 Choosing an anti-aliasing filter

A sensible procedure for determining a sampling rate starts with the person who wants the data (say, the scientist) stating clearly the highest frequency, $f_0$, that must be resolved so that the data are useful for their intended purpose. Call $f_0$ the "band of interest". The extent of this band is a scientific question. Designers, engineers and other suppliers have no say in its choice beyond telling you how much it will cost. It is then the responsibility of the suppliers to produce a measurement system that will collect data, throughout the band of interest, with no more than some $X$ amount of aliasing. Unfortunately, there is no completely objective criterion for specifying $X$, because its value dependends on the characteristics of the signal that you wish to measure and on the resolution of your measurement system (see sections 3 and 4). RSI usually specifies a maximum aliasing of $X \approx 1 \times 10^{-5}$ in the band of interest ($f < f_0$), using a cascade of two $4^{\text{th}}$-order Butterworth filters.

There are at least two design methods; conservative and minimum sampling rate.

## 6.1 Conservative design

To attain an attenuation such as $X = 1 \times 10^{-5}$, one simply finds the frequency at which the filter achieves this value. For the $8^{\text{th}}$-order Butterworth filter of Figure 4, this occurs at a frequency $f_X = 4.3 f_0$. The sampling rate is then $f_s = 2 f_X = 8.6 f_0$. Aliasing at all frequencies is then guaranteed to be smaller than $X$. However, the sampling rate is unnecessarily high, wasting storage space and data transmission bandwidth.

Although it is quite possible to design and build $8^{\text{th}}$-order Butterworth filters, a cascade of two $4^{\text{th}}$-order Butterworth filters is more practical because, for lower orders, the filters are more stable and components are more abundant, and the results are nearly as effective as those of a single $8^{\text{th}}$-order filter.

## 6.2 Minimum sampling rate design

A continuous-domain signal with a frequency $f_N < f < 2f_N$ is aliased to a frequency $f_a = 2f_N - f$. The frequency folds back around the Nyquist frequency. In general, the continuous-domain frequency, $f$, is aliased to the frequency

$$\begin{aligned} f_a &= (-1)^n \left(f - k_n f_N\right), \\ k_n &= n + \frac{1-(-1)^n}{2}, \quad \text{and} \\ n &= 0,\ 1,\ 2,\ .... \end{aligned} \tag{12}$$

where $n = 0$ represents the non-aliased Nyquist band, $n = 1$ represents the first aliased branch, and so on. The effect of aliasing can be illustrated by plotting the response of the anti-aliasing filter as it appears in the sampled data (Figure 5, red). The first branch ($n =$

1) is the most important one because it has the least amount of attenuation. Aliasing by way of higher branches requires a large continuous-domain signal in order for it to bias (or contaminate) the band of interest of 0 to $f_0$. Choosing a sampling rate of $f_s = 5.7f_0$ makes the first branch of the alias smaller than $1 \times 10^{-5}$ in the entire band of interest. For this cascade of two 4$^{\text{th}}$-order Butterworth filters, the conservative design would have used a sampling rate of $2 \times 4.7 = 9.4\,\text{Hz}$. The rate of the minimum sampling rate design is 1.65 times smaller.

Again, the specifications of section 4 are very conservative. The design used by RSI has $f_s = 5.22f_0$ (Figures 6 and 7), for technical reasons. The band of interest for typical instruments is $f_0 \approx 100\,\text{Hz}$ and so the anti-aliasing filters are set to $f_0 = 98\,\text{Hz}$ and the sampling rate is set to $f_s = 512\,\text{s}^{-1}$.

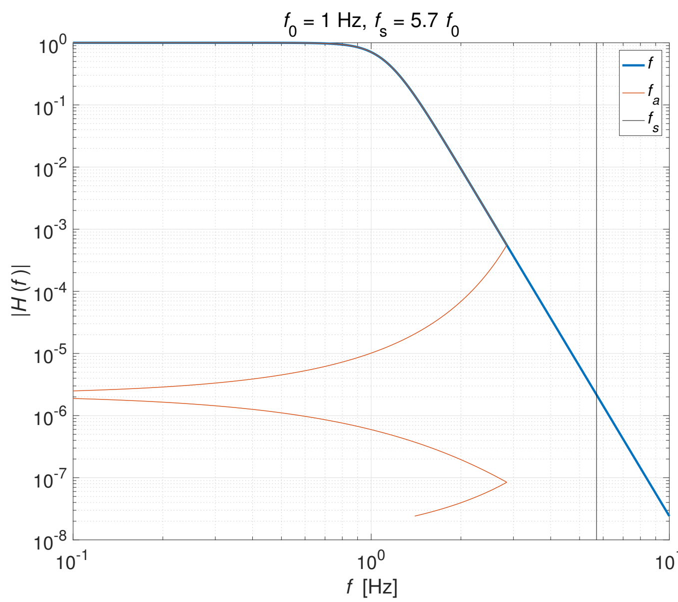


Figure 5: The response magnitude of a cascade of two 4$^{\text{th}}$-order Butterworth filters set for a combined half-power response at $f_0 = 1\,\text{Hz}$. Blue – response in the continuous domain. Red – response after sampling at $f_s = 5.7\,\text{Hz}$. Black – 5.7 Hz.

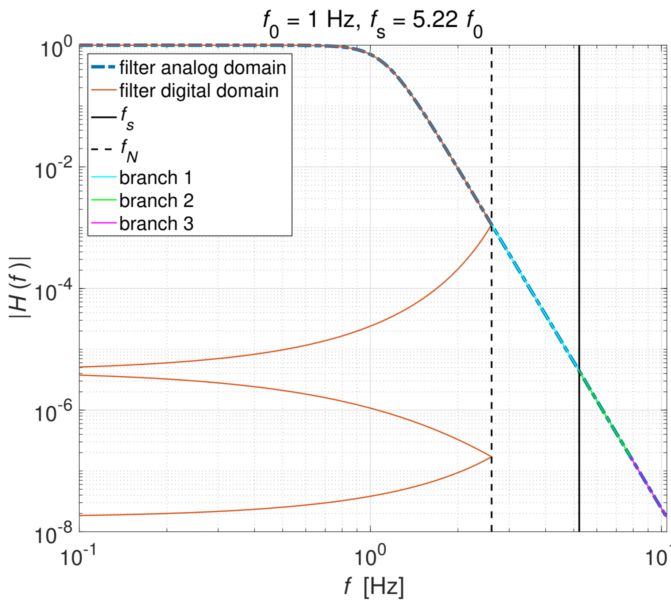


Figure 6: Same as Figure 5, but with the RSI implementation of $f_s = 5.22\,\mathrm{Hz}$. The first three branches of the aliased signals have been highlighted in cyan, green and magenta.

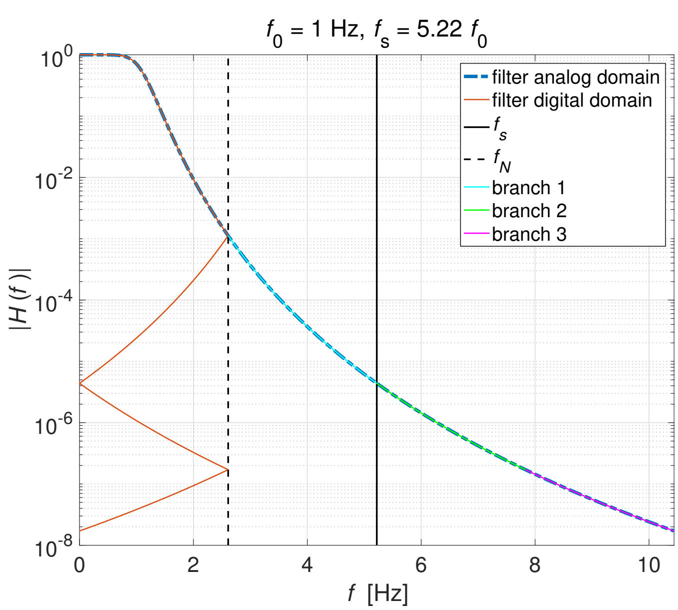


Figure 7: Same as Figure 6, but with a linear frequency axis.

# 7 Identifying an anti-aliasing filter

You may not know if the system that was used to collect your data contained an anti-aliasing filter. Fortunately, there will be evidence of an anti-aliasing filter in your spectral estimates, if such a filter was used. The rapid decrease of a signal with increasing frequency (which is a hallmark of a good anti-aliasing filters) is quite unnatural and should be visible in the spectrum of your signal. That is, you should see a rapid decrease of your spectrum, with increasing frequency, near the Nyquist frequency (Figure 3). If you do not see this characteristic, then your data may be aliased. If you also know the full-scale range and the number of bits, $B$, of the sampler, then you can use Figure 3 to estimate the quality of the sampler.

————End of Document————