Title

Analysis and Design of Digital Parametric Filters Without Frequency Warping

ISBN

9798293867165

Author

Muranov, Mark Samuel

Publication Date

2025-09-23

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Analysis and Design of Digital Parametric Filters Without Frequency Warping

A Thesis submitted in partial satisfaction of the
requirements for the degree Master of Science

in

Electrical Engineering (Signal and Image Processing)

by

Mark Muranov

Committee in charge:

Professor Fred Harris, Chair
Professor Dinesh Bharadia, Co-Chair
Professor Truong Nguyen

2025

The Thesis of Mark Muranov is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2025

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Fred Harris for his support as the chair of my committee. His feedback and research guidance proved invaluable.

I would also like to thank Professors Nguyen and Bharadia for their excellent digital signal processing courses. The courses were instrumental in my education.

ABSTRACT OF THE THESIS

Analysis and Design of Digital Parametric Filters Without Frequency Warping

by

Mark Muranov

Master of Science in Electrical Engineering (Signal and Image Processing)

University of California San Diego, 2025

Professor Fred Harris, Chair
Professor Dinesh Bharadia, Co-Chair

Filters are essential tools in audio engineering. For decades, music producers, DJs, and mixing engineers have relied on a standard set of parametric filters to shape audio signals. These filters—low-pass, high-pass, band-pass, notch, peaking, low-shelf, and high-shelf filters—are typically implemented as biquadratic structures with two poles and two zeros. Their parametric design allows users to intuitively adjust key characteristics in real time.

With the transition to digital systems, the bilinear transform (BLT) became a common method for adapting analog filter designs to the discrete domain. However, this conversion introduces frequency warping, distorting the intended frequency responses.

This thesis explores existing solutions to this problem and introduces a novel approach. We examine the BLT, brute-force oversampling, two prior papers which correct for frequency warping, and finally propose a new technique that integrates and improves upon these methods to reduce frequency distortion in digital filter implementations.

Chapter 1

Introduction

1.1 Motivation

Parametric filters have become an industry-standard tool across nearly all areas of audio and music production. Mixing engineers use them to enhance clarity and separation in mixes, while music producers rely on them as creative tools for sound design. There exist 7 commonly used parametric filter types: low-pass, high-pass, band-pass, notch, peaking, low-shelf, and high-shelf [2]. Each of these filters can be implemented as a second-order system (biquad) in the continuous-time Laplace domain, with two poles and two zeros.

Parametric filters are typically controlled using three parameters: cutoff frequency ω_0, quality factor Q, and, in some cases, gain g. The cutoff frequency defines either the center or edge of the filter's frequency range. The quality factor is inversely proportional to the damping ratio ζ [4] and determines the filter's bandwidth or resonance. The gain parameter, present in filters like the peaking and shelving types, specifies the amount of amplification or attenuation at the targeted frequencies. These parameters are intuitive to audio engineers and allow for straightforward, deterministic computation of filter coefficients, enabling fast and flexible tuning.

When analog filters were adapted to the digital domain, the bilinear transform (BLT) was widely adopted due to its ability to preserve stability, causality, and phase characteristics [1][3][12]. However, the BLT introduces a significant drawback: frequency warping. This occurs because the BLT maps the infinite frequency range of the continuous-time domain onto the

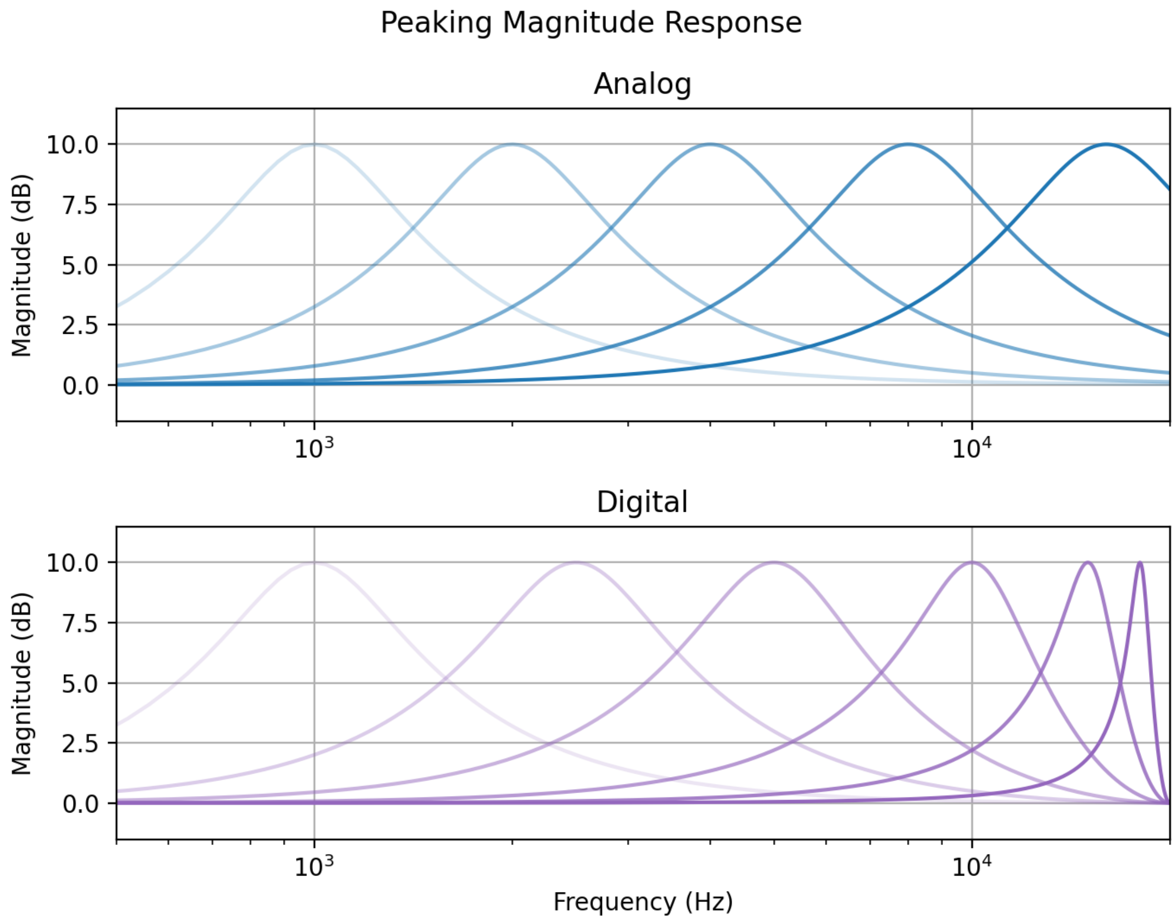

Figure 1.1. Demonstration of frequency cramping of the digital peaking filter across different values of f_0 ($Q = 1$, $g = 10$dB, $F_s = 40$kHz).

finite range of the digital (discrete-time) domain in a nonlinear fashion [1]. This results in a non-uniform distribution of frequencies, particularly as they approach the Nyquist limit, causing the filter response to appear "squished" or "cramped." This phenomenon is commonly referred to as *frequency warping* or *digital cramping* and is illustrated in Figure 1.1.

As a result, the behavior of the digital filters deviates from their analog counterparts, with the response varying significantly depending on the cutoff frequency and the system's sampling rate. This inconsistency not only complicates filter design but also results in audible differences when the same filter is used at different sampling rates—an undesirable outcome in modern digital audio workflows.

The goal of this thesis is to investigate solutions to the frequency warping problem in digital parametric filters. We will compare existing approaches in terms of accuracy, implementation complexity, and computational cost, and propose a novel hybrid method that combines the strengths of previous techniques.

1.2 The Relevance of Phase in Audio

In digital communications, linear phase finite impulse response (FIR) filters are commonplace, due to their ability to preserve symbols; the lack of non-linear phase distortion preserves the wave-shape. However, linear phase filters incur increased latency and computational costs.

On the other hand, linear phase filters tend to be unnecessary in audio. In particular, it has been shown that people have difficulty perceiving changes to phase caused by filters [15]. However, people are sensitive to latency caused by linear phase filters. As a result, minimum phase infinite impulse response (IIR) filters are desirable for their low latency and minimal computational costs. In fact, all the filters presented in the Audio EQ Cookbook [2] are minimum phase.

A real-world demonstration of this can be found in hearing aids. When initially conceived, digital hearing aids utilized linear-phase filters. However, over the years, these devices transitioned to using minimum-phase filters. Due to this change, the hearing aids' costs and latency were reduced in exchange for no perceivable downside [13].

Due to these benefits, our attention will be solely focused on minimum phase IIR filters. In addition, we will prioritize matching the magnitude responses over the phase responses, as differences in phase will be difficult to perceive.

Chapter 2

Existing Techniques

2.1 Bilinear Transform

The most commonly used technique for adapting the parametric analog filters is the bilinear transform (BLT). Audio software, ranging from digital audio workstations (DAWs) like FL Studio and Ableton Live to internet browsers like Google Chrome (used in the Web Audio API) [2] use the BLT-derived filters.

The Audio EQ Cookbook [2] describes the original analog prototype filters and the corresponding BLT-derived digital filter coefficients. Applying the BLT is a straightforward process: substitute s in the analog filter transfer function with the following term:

$$s \leftarrow c \cdot \frac{1 - z^{-1}}{1 + z^{-1}} \tag{2.1}$$

Where c is the bilinear constant.

As mentioned in the introduction, the BLT maps the infinite analog plane onto the finite and periodic digital plane. To achieve this, a non-linear mapping is used [1]. The BLT frequency mapping is described by Equation 2.2.

$$f_d = \tan\left(\frac{\pi f_c}{F_s}\right) \cdot \frac{F_s}{\pi} \tag{2.2}$$

Where f_c is the original analog frequency, f_d is the warped digital frequency, and F_s is

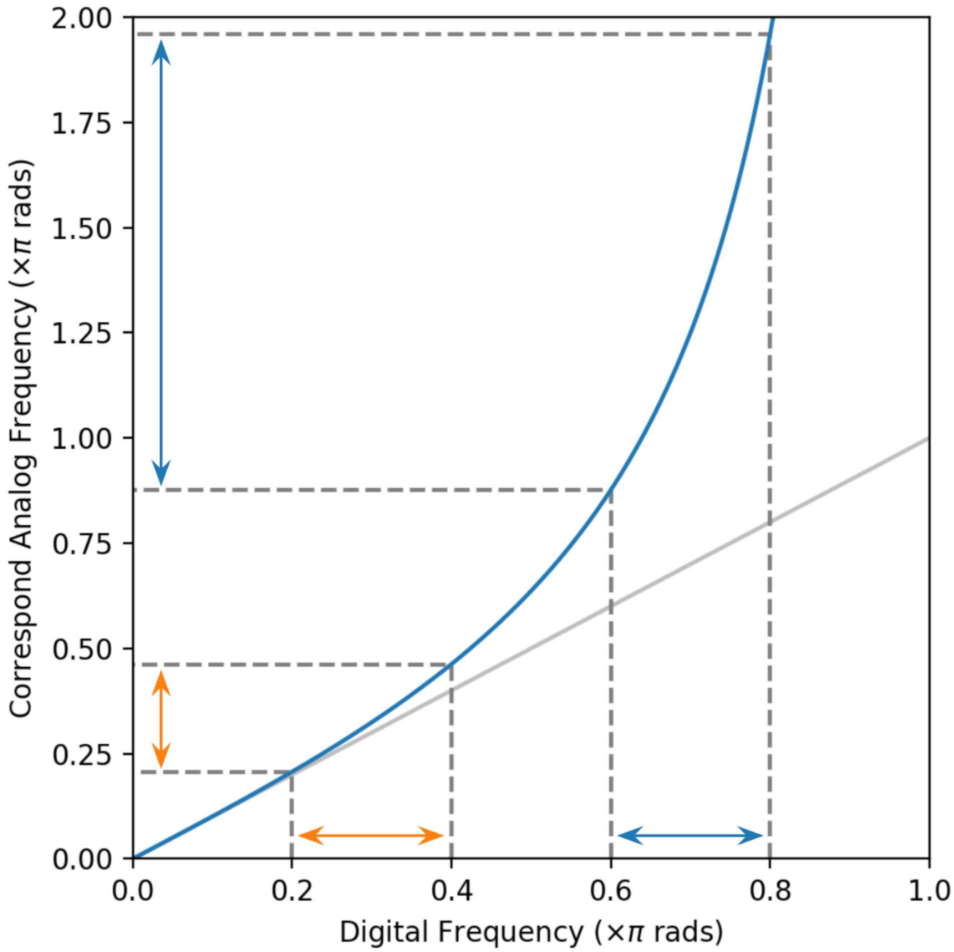

Figure 2.1. Bilinear transform frequency mapping.

the sampling rate.

The tangent function behaves approximately linearly near $f_d = f_c = 0$ but increases rapidly as the digital frequency f_d approaches the Nyquist frequency $f_d = \frac{F_s}{2}$. In fact, as the digital frequency approaches the Nyquist frequency $f_d \rightarrow \frac{F_s}{2}$, the continuous frequency approaches infinity $f_c \rightarrow \infty$. Consequently, higher analog frequencies are increasingly compressed as they are mapped to the digital domain, as demonstrated in Figure 2.1.

To help mitigate this issue, a technique known as "pre-warping" is applied. Typically, the bilinear constant is used to pre-warp a specific frequency. For instance, a commonly used constant is shown in Equation 2.3.

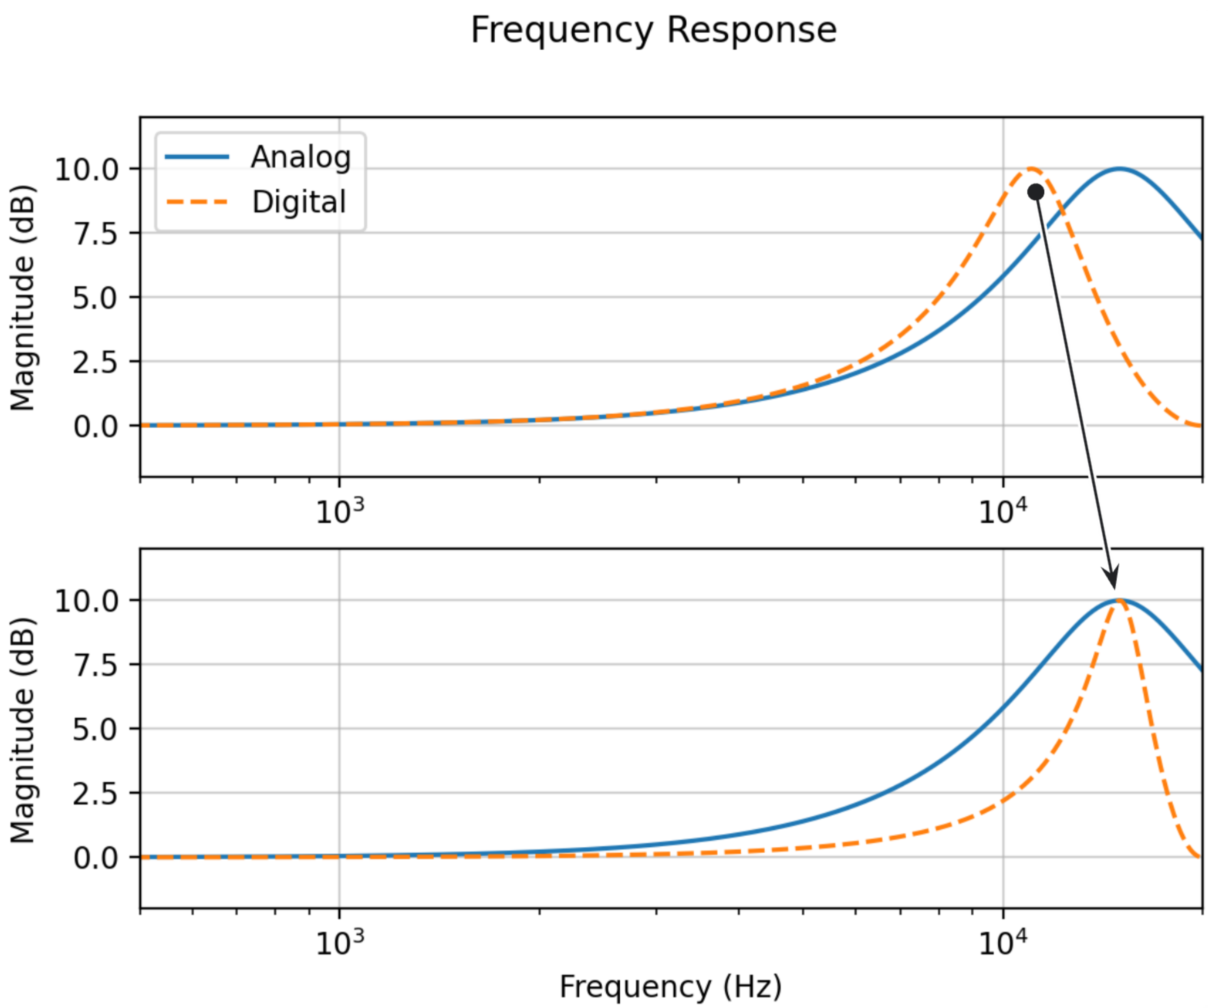

Figure 2.2. Parametric peaking filter without pre-warping (top) and with pre-warping (bottom). $f_0 = 15\text{kHz}$, $Q = 1$, $g = 10\text{dB}$, $F_s = 40\text{kHz}$.

$$c = \frac{\omega_0}{\tan(\omega_0/2)} \tag{2.3}$$

This constant will ensure that the cutoff frequency ω_0 of the filter is pre-warped such that its position is preserved between the analog and digital versions. However, as demonstrated in Figure 2.2, all neighboring frequencies are still warped by the BLT.

2.1.1 Deriving the Digital Peaking Filter

As an example of using the BLT, we will derive the digital peaking filter as found in the Audio EQ Cookbook [2]. The analog peaking filter has the following transfer function:

$$H(s) = \frac{s^2 + s\frac{g\cdot\omega_0}{Q} + \omega_0^2}{s^2 + s\frac{\omega_0}{Q} + \omega_0^2} \tag{2.4}$$

Notice that this transfer function is different from the one described in the Audio EQ Cookbook [2]. If we substitute $Q \leftarrow Q\sqrt{g}$ into the above transfer function, it will match the Audio EQ Cookbook. Having only a single gain g term simplifies the math and will be used throughout this thesis.

Applying the BLT onto this transfer function and collecting like terms results in the following digital transfer function:

$$H(z) = \frac{z^{-2}(n^2 - n\frac{g}{Q} + 1) + z^{-1}(-2n^2 + 2) + (n^2 + n\frac{g}{Q} + 1)}{z^{-2}(n^2 - n\frac{1}{Q} + 1) + z^{-1}(-2n^2 + 2) + (n^2 + n\frac{1}{Q} + 1)} \tag{2.5}$$

Where $n = 1/\tan(\omega_0/2)$. From this result, we can extract the digital coefficients:

$$\begin{aligned} b_0 &= n^2 + n\tfrac{g}{Q} + 1, \quad b_1 = 2(1 - n^2), \quad b_2 = n^2 - n\tfrac{g}{Q} + 1, \\ a_0 &= n^2 + n\tfrac{1}{Q} + 1, \quad a_1 = 2(1 - n^2), \quad a_2 = n^2 - n\tfrac{1}{Q} + 1 \end{aligned}$$

The filter can then be implemented with Direct Form 1 [12] with the difference equation in Equation 2.6.

$$y[n] = \left(\frac{b_0}{a_0}\right)x[n] + \left(\frac{b_1}{a_0}\right)x[n-1] + \left(\frac{b_2}{a_0}\right)x[n-2] - \left(\frac{a_1}{a_0}\right)y[n-1] - \left(\frac{a_2}{a_0}\right)y[n-2] \tag{2.6}$$

Or more efficiently with Direct Form 2 [12]. It is common to divide all the coefficients by a_0 when initially computing the coefficients, to avoid needing any divide operations per sample.

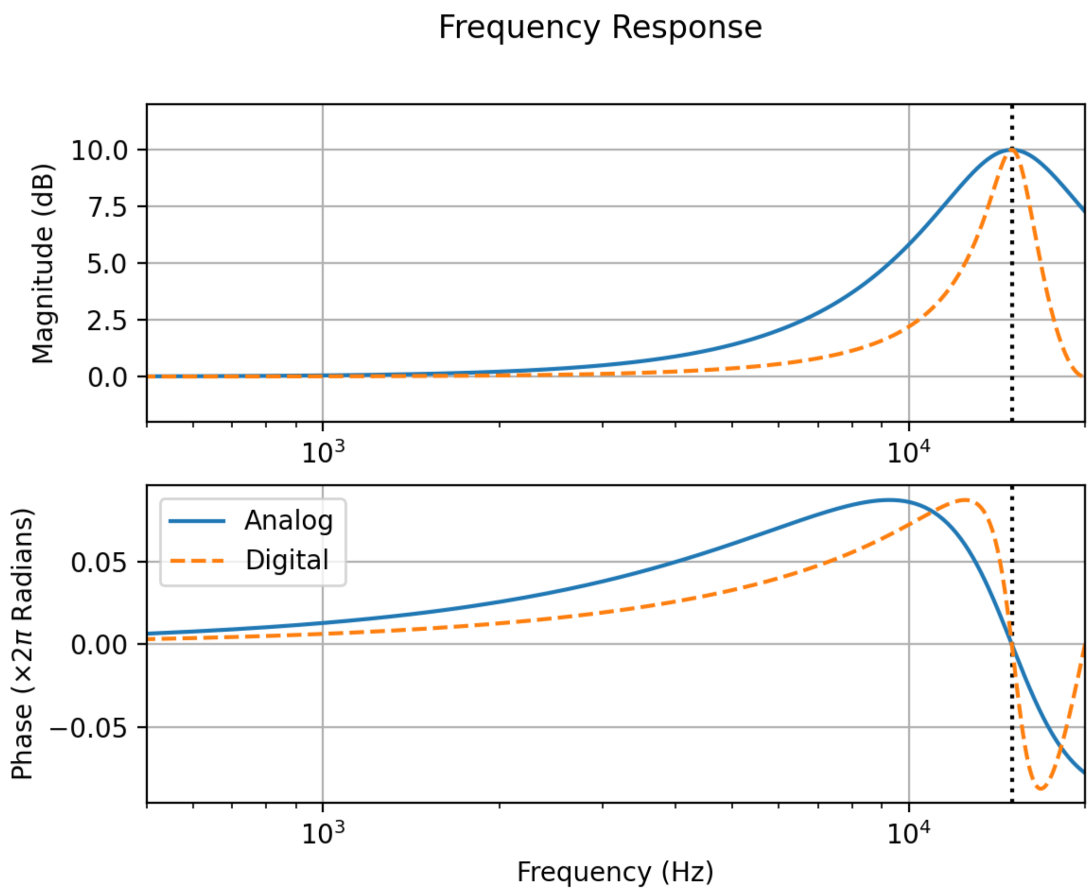

Figure 2.3. Analog and BLT-derived digital peaking filter. $f_0 = 15\text{kHz}$, $Q = 1$, $g = 10\text{dB}$, $F_s = 40\text{kHz}$.

2.1.2 Analysis

The warping is most noticeable when setting the cutoff frequency to a value near Nyquist, for instance, when the cutoff frequency f_0 is three quarters that of the Nyquist frequency $F_s/2$. The analog and digital parametric peaking filter responses are displayed in Figure 2.3.

As predicted, while the filter has the correct cutoff frequency and gain, the neighboring frequencies display large differences between the two domains.

2.2 Oversampling

The first approach we will be exploring is oversampling. We previously observed that the frequency warping is most prevalent for frequencies approaching the Nyquist frequency. By increasing the sampling-rate via up-sampling/interpolation, the Nyquist frequency is increased. For instance, up-sampling by 2x also doubles the Nyquist frequency. And because we observed that the non-linear mapping mostly exists in the upper-half of the bandwidth, up-sampling by 2x moves most of the warping to above the original Nyquist limit. The frequency warping is mostly eliminated throughout the original bandwidth and the signal can be safely down-sampled/decimated back to the original sampling-rate.

However, this approach is not without its downsides. Most prominently, the effect of 2x oversampling means the filters must operate on twice as many samples, doubling processing costs. In addition, an interpolation filter must be used, further increasing computational complexity. The interpolation filter also imparts its own frequency response and latency.

Let us take a look at a specific implementation. We will use the peaking filter as presented in the Audio EQ Cookbook [2] and the half-band low-pass polyphase filter described in Valenzuela's and Constantinides' paper [14].

The latter paper provides up-sampling and down-sampling filters whose coefficients can be computed via an analytic design method. We only need to specify two design constraints: the normalized transition-band width ω_t (ranging between 0 and 0.5) and the maximum stop-band attenuation g_s. Choosing $\omega_t = 0.1$ and $g_s = -60\text{dB}$ results in a 9th-order two-path filter, consisting of two all-pass filters per path. This structure is shown in Figure 2.4.

As outlined in the paper, this filter offers excellent filtering performance at a very low computational cost and is employed in the widely-used JUCE audio processing framework. It comfortably satisfies our design constraints, as demonstrated in Figure 2.5. Based on this, our processing chain for the oversampling approach is as follows (also shown in Figure 2.6):

1. Up-sample/interpolate the signal by 2x using the interpolation filter $H_{\uparrow}(z)$.

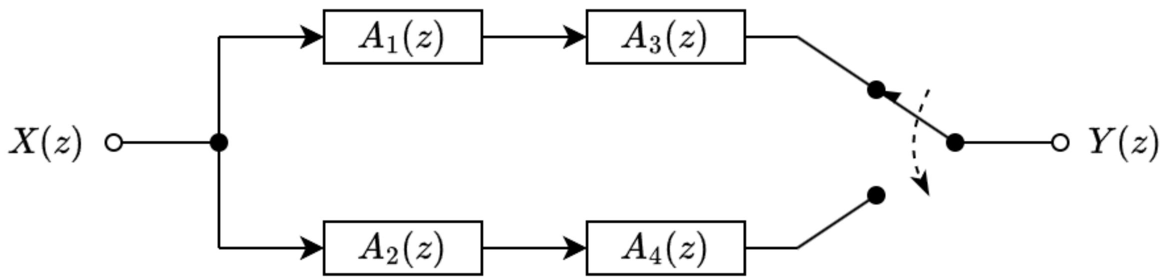

Figure 2.4. Interpolation filter structure $H_{\uparrow 2}(z)$.

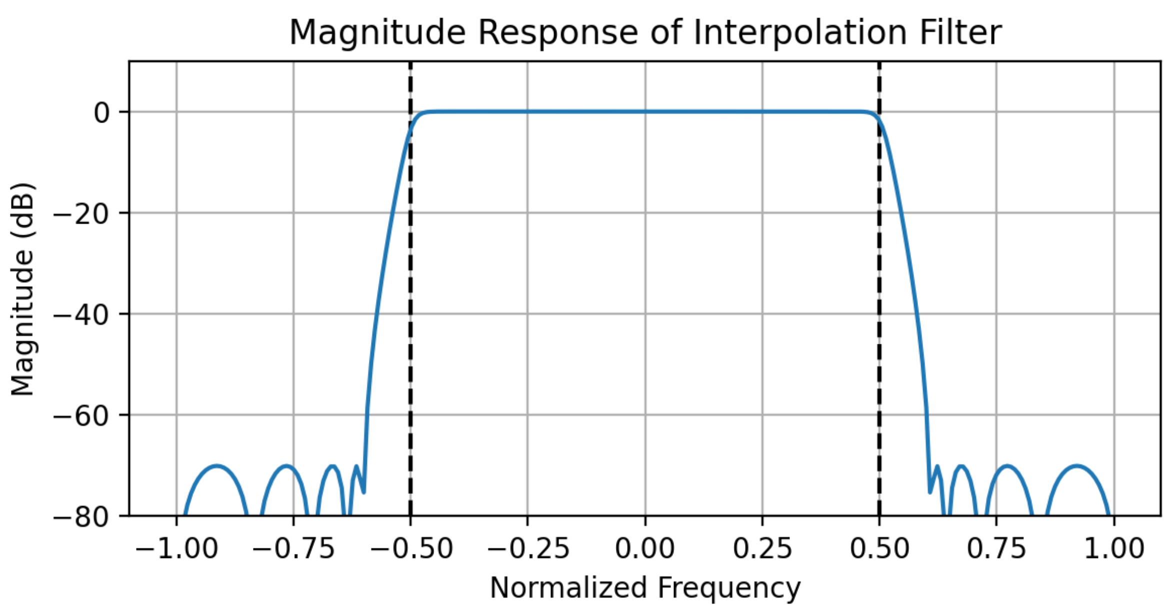

Figure 2.5. Interpolation filter response ($\omega_t = 0.1$, $g_s = -60\text{dB}$)

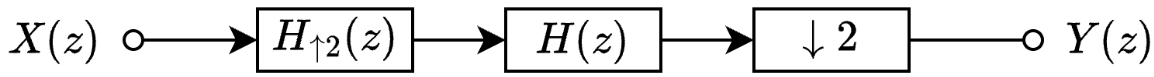

Figure 2.6. Oversampling system block diagram.

2. Process the parametric filter $H(z)$ at the increased sample-rate.

3. Down-sample/decimate by 2x by discarding every even-indexed sample.

Note that we do not apply a decimation filter during the down-sampling step because the up-sampling process has already attenuated frequencies above the original Nyquist limit. Since the processing we perform at the higher sample rate is linear and time-invariant (LTI), it does not introduce any new harmonics that could otherwise cause aliasing [6]. To verify this behavior, we analyze the system's impulse response, as shown in Figure 2.7.

The system response highlights the improved frequency response achieved by the oversampled approach compared to the bilinear transform (BLT)-derived filter. While not a perfect match, the magnitude response more closely resembles that of the original analog filter.

However, despite this improvement, the oversampled approach presents two key drawbacks. First is the added computational complexity. Because the parametric filters operate at the increased sample rate, the number of multiplies per sample effectively doubles. This cost is further compounded by the fact that users often employ multiple filters in series. Moreover, this does not account for the additional overhead introduced by the interpolation filter. Second, the interpolation process imparts its own phase response onto the signal, introducing latency and phase distortion—artifacts that are not present in the analog counterpart.

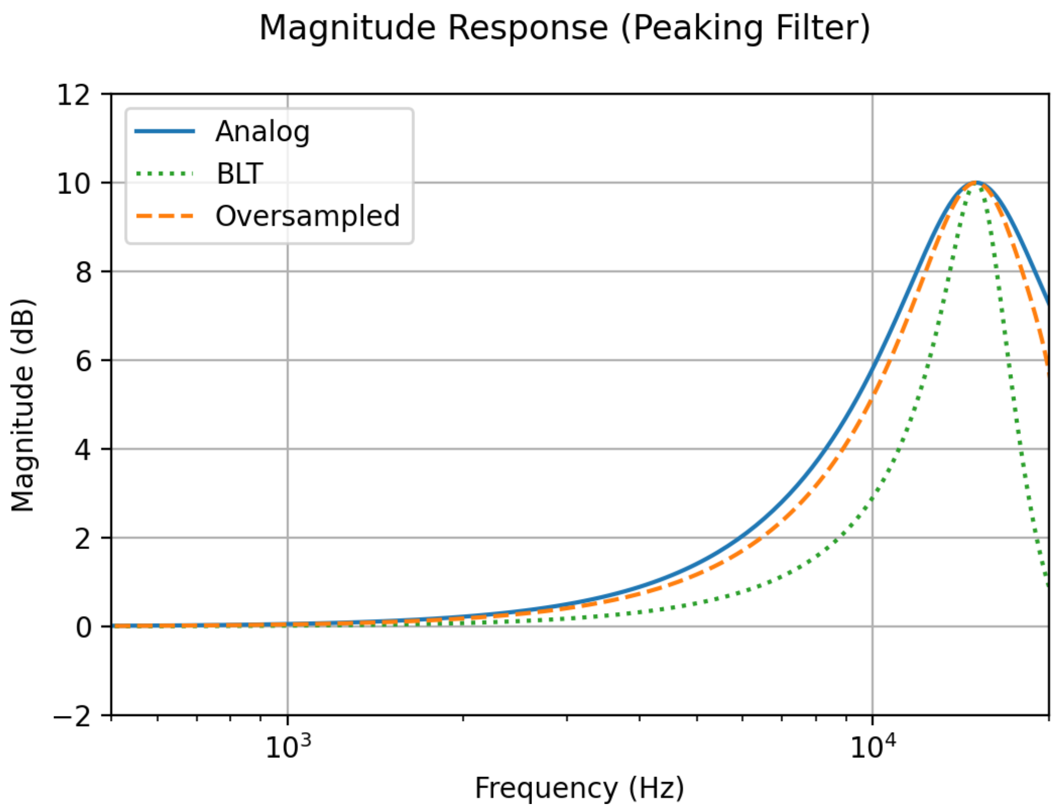

Figure 2.7. Peaking filter parameters: $F_s = 44.1\text{kHz}$, $f_0 = 15\text{kHz}$, $Q = 1$, $g = 10\text{dB}$

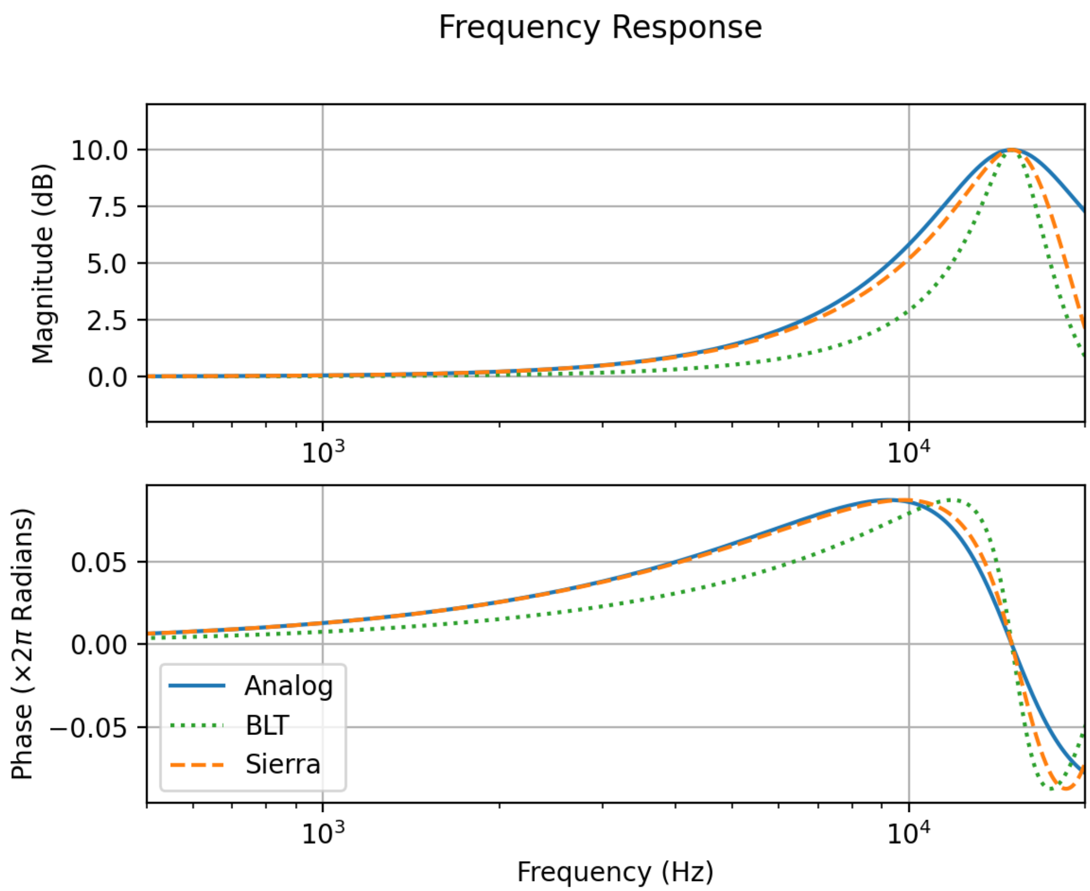

Figure 2.8. Comparison with Sierra's Peaking filter. $F_s = 44.1$kHz, $f_0 = 15$kHz, $Q = 1$, $g = 10$dB

2.3 Sierra's Method

While Sierra does not intend to directly address the frequency warping issue, the proposed technique does mitigate its effects [11]; Sierra's paper focuses on the design of parametric peaking filters that support cutoff frequencies above the Nyquist limit. Since our focus is on frequency warping, we will only briefly discuss the full details of Sierra's solution.

Let's observe how the frequency warping effect impacts the magnitude response of the parametric peaking filter. As demonstrated in Figure 2.8, the warping appears to reduce the bandwidth of the peak. Coincidentally, the quality Q parameter also controls the peak bandwidth. So, what if we were to preemptively lower the value of Q to increase the bandwidth in anticipation of the warping effect? This pre-warped quality Q_d should be a function of Q.

Logically, for cutoff frequencies f_0 near DC, we would want $Q_d = Q$, since the warping effect is minimally present. But, as the cutoff frequency approaches Nyquist, Q_d should decrease, approaching $Q_d = 0$. Sierra proposes the following simple Q_d formula in Equation 2.7, where $c = \omega_0/\tan(\omega_0/2)$ is the bilinear constant as found in Equation 2.3.

$$Q_d = Q \cdot \frac{c}{2} \tag{2.7}$$

For small values of the normalized cutoff frequency ω_0, we can utilize the small angle approximation of the tan function. In particular, the bilinear constant is approximately $c \approx 2$ as $\omega_0 \to 0$.

This simple tweak to the Q parameter provides a significant improvement to the magnitude response, as seen in Figure . This can be a great quick fix for the peaking filter, as it maintains all the beneficial properties of the standard BLT-derived filters, such as inexpensive computation, stability, and phase properties. However, even with this adjustment, the digital magnitude response still significantly deviates from its analog counterpart, as seen in Figure 2.8. Furthermore, approach is only suitable for filters in which the quality Q parameter adjusts the bandwidth, which only encompasses the peaking, band-pass, and notch filters.

As mentioned, Sierra's paper proposes a peaking filter design which allows cutoff frequencies above the Nyquist limit. Recall that the BLT maps the entirety of the infinite analog frequency plane to the finite and periodic digital plane. As a result, going above the Nyquist limit results in undefined behavior. In particular, let's take a look at the bilinear constant $c = \omega_0/\tan(\omega_0/2)$. Due to the tan function, the bilinear constant contains discontinuities at multiples of $\omega_0 = \pi$, which correspond to the Nyquist limit. Sierra describes how this discontinuity results in unstable filters with distorted frequency responses [11].

To combat these issues, Sierra proposes a new bilinear constant in Equation 2.8. The new constant closely matches the original for small values of ω_0, but remains positive even as $\omega_0 \to \infty$.

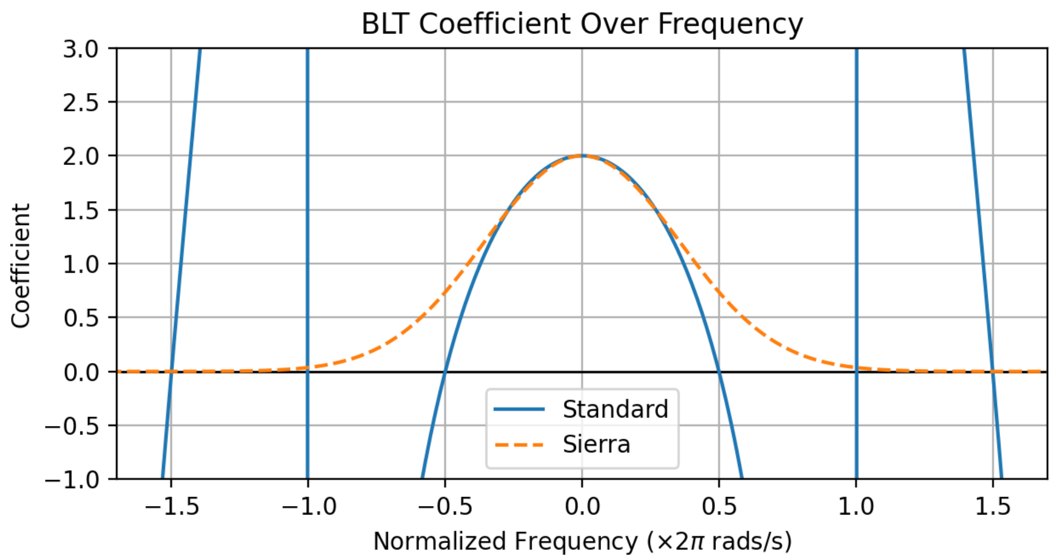

Figure 2.9. Comparison of the standard BLT coefficient to the Sierra's proposed coefficient.

$$c = 2e^{-(\frac{\omega_0}{\pi})^2} \tag{2.8}$$

As displayed in Figure 2.9, the proposed constant closely matches the original when $\omega_0 << 2\pi F_s$, but as $\omega_0 \to 2\pi F_s$, the new constant increases at a slower rate. In fact, as $\omega_0 \to \infty$, $c \to 0$, meaning the constant will never be negative.

Unlike the standard bilinear constant, this proposed constant does not correctly pre-warp the cutoff frequency ω_0. To correct for this, we must match the gain at the distorted ω_0 position, which we will call ω_d, in addition to computing the pre-warped quality Q_d.

To find ω_d, we can use the mapping described in Equation 2.2 in conjunction with the proposed bilinear constant in Equation 2.8 as follows:

$$\omega_d = 2\tan^{-1}(\omega_0/c) \tag{2.9}$$

This process actually results in the cutoff frequency to be in a slightly incorrect position. Despite this, we can closely match the analog magnitude response by adjusting the peak gain. In the case of the parametric peaking filter, we can find the magnitude of the transfer function in Equation 2.4 by taking its norm, $||H(\omega)||$. The goal is to find the analog filter gain at frequency ω_d. We can do this by substituting $\omega \leftarrow \omega_d$ in the magnitude function $||H(\omega)||$:

$$g_d^2 = \frac{(\omega_d^2 - \omega_0^2)^2 + g^2\omega_d^2\omega_0^2/Q^2}{(\omega_d^2 - \omega_0^2)^2 + \omega_d^2\omega_0^2/Q^2} \tag{2.10}$$

To compute the corrected quality factor Q_d, Sierra uses the following equation:

$$Q_d = Q \cdot \frac{c}{2F_s} \cdot \frac{g}{g_d} \tag{2.11}$$

Substituting these corrected parameters in the peaking filter transfer function in Equation 2.4 produces:

$$H_d(s) = \frac{s^2 + s\frac{g_d \cdot \omega_d}{Q_d} + \omega_d^2}{s^2 + s\frac{\omega_d}{Q_d} + \omega_d^2} \tag{2.12}$$

Taking the bilinear transform of $H_d(s)$ with the modified bilinear constant produces the desired digital parametric peaking filter. Comparing the frequency response in Figure 2.10, Sierra's method provides significant improvements over the BLT-derived filter. In particular, for frequencies below ω_d, both the magnitude and phase response are very similar to the analog prototype. However, above ω_d, the response deviates significantly, being more similar to the standard BLT filter. When using a traditional sample-rate, such as 44.1kHz or above, the deviating region approaches the upper limit of human hearing [9], mitigating the perceived warping.

This approach improves upon the standard BLT while preserving order, stability, and remaining minimum phase. As a result, no additional latency or computational complexity is introduced. The only caveat is the increased number of operations needed to compute the coefficients.

While Sierra's paper only covers the parametric peaking filter, there exists a generalization

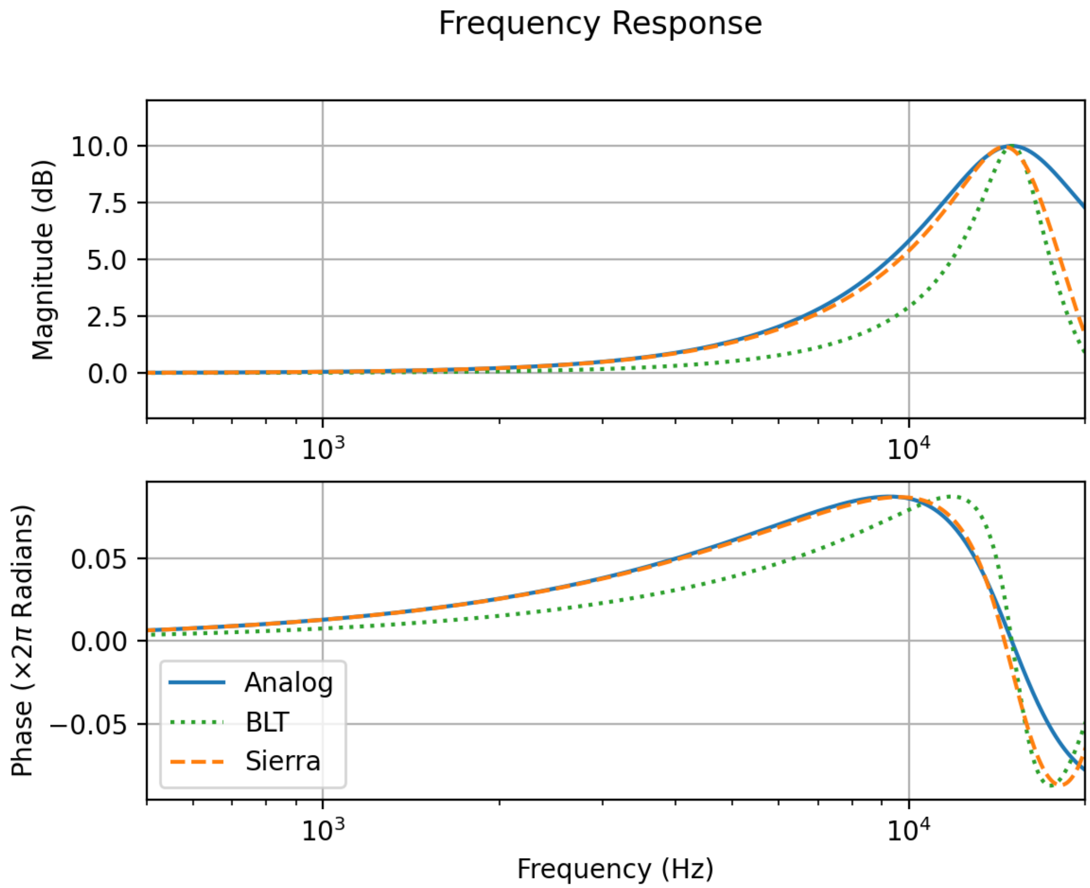

Figure 2.10. Comparison with Sierra's Peaking filter (w/ proposed bilinear constant). $F_s =$ 44.1kHz, $f_0 =$ 15kHz, $Q = 1$, $g =$ 10dB

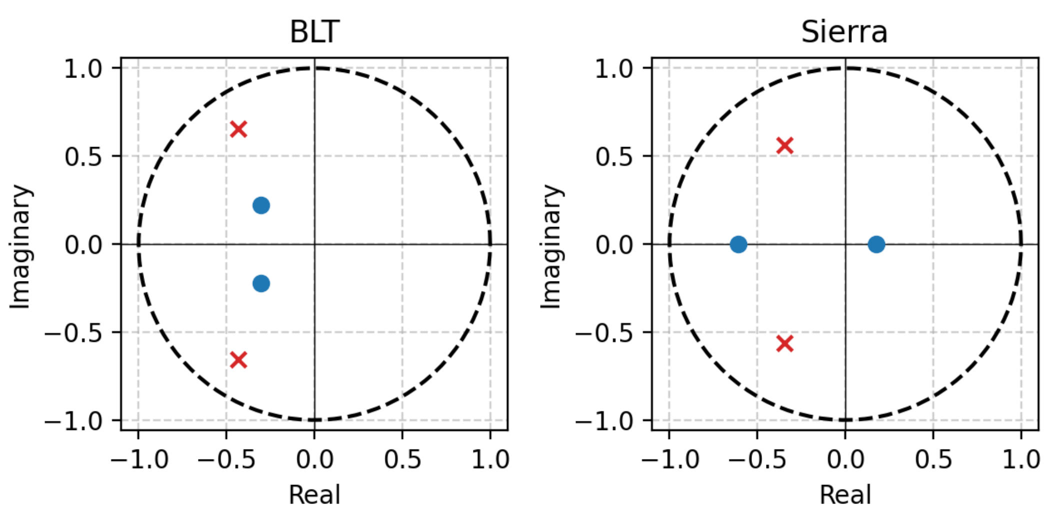

Figure 2.11. Pole-zero diagrams ($F_s =$ 44.1kHz, $f_0 =$ 15kHz, $Q = 1$, $g =$ 10dB).

[4] which supposedly allows Sierra's approach to be used for any second-order system. Whether this approach can be used for the other parametric filter types remains to be seen.

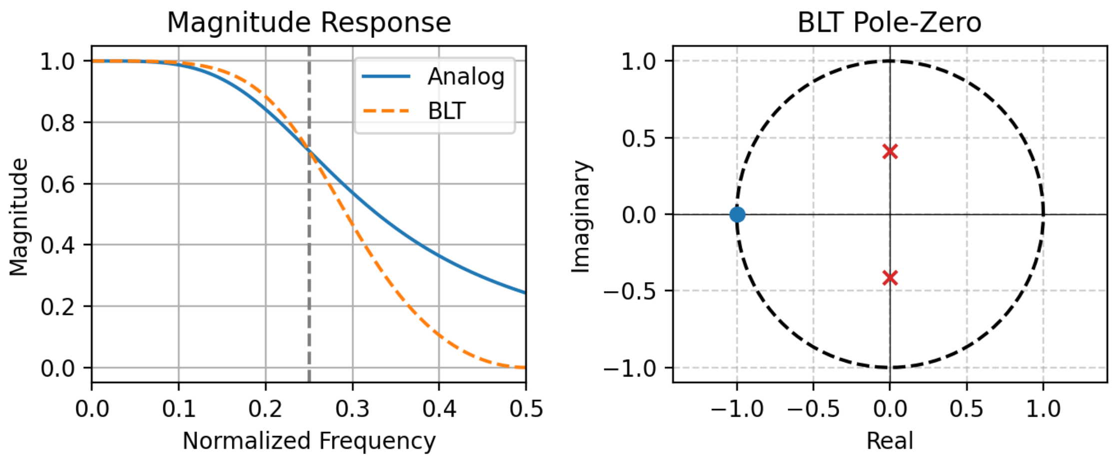

Figure 2.12. BLT-derived low-pass filter analysis ($f_0 = F_s/4$, $Q = 1/\sqrt{2}$).

2.4 Orfanidis' Method

Orfanidis addresses the frequency warping issue of the BLT with a method he calls ”Prescribed Nyquist-Frequency Gain” [7]. This approach stems off of one key observation: the BLT maps the analog filter response at $s \to \infty$ onto the digital filter response at $z = -1$. The issue with this mapping is best demonstrated with the parametric low-pass filter [2]:

$$H(s) = \frac{\omega_0^2}{s^2 + s\frac{\omega_0}{Q} + \omega_0^2} \tag{2.13}$$

Notice what happens as we take the limit $s \to \infty$. Since the denominator contains s^2 while the numerator is constant, $H(s)$ approaches 0. We can conclude that $H(s)$ contains two zeros at $s \to \infty$. As discussed earlier, these two zeros will be mapped onto $z = -1$. We can verify this by taking the BLT and analyzing the pole-zero diagram, as seen in Figure 2.12.

How does this impact the frequency response? For the analog low-pass filter, the magnitude response approaches 0 as $\omega \to \infty$ but never quite reaches it. On the other hand, for the digital filter, the magnitude response *does* reach 0 at the Nyquist frequency. As a result, the transition band of the analog filter remains infinitely large, regardless of cutoff frequency, which cannot be said about the digital filter. In fact, as the cutoff frequency f_0 approaches the

Nyquist frequency $F_s/2$, the digital filter's transition band shrinks, further deviating from the analog filter's slope (Figure 2.12).

This same observation can be applied to any BLT-derived filter, including all the second-order parametric filters. Is there anything we can do to mitigate this issue? Let's analyze Orfanidis' approach. First, he constructs a general second-order filter as follows:

$$H_d(s) = \frac{G_1 s^2 + Bs + G_0 W^2}{s^2 + As + W^2} \tag{2.14}$$

Here we have five parameters, G_0, G_1, A, B, and W we can tweak. Important to note are the G_0 and G_1 parameters. As we did with the low-pass example, let's observe what happens to $H_d(s)$ as $s \to \infty$. The two s^2 terms grow the fastest, so the others can be ignored. As a result, the transfer function simplifies to the following:

$$\lim_{s \to \infty} H_d(s) = \lim_{s \to \infty} \frac{G_1 s^2}{s^2} = G_1 \tag{2.15}$$

Similarly, let's see what happens at DC (when $s = 0$):

$$H_d(s = 0) = \frac{G_1 \cdot 0 + B \cdot 0 + G_0 W^2}{0 + A \cdot 0 + W^2} = \frac{G_0 W^2}{W^2} = G_0 \tag{2.16}$$

The G_0 and G_1 terms control the gain at DC and at the Nyquist frequency, respectively. Going back to the low-pass example, this would allow us to design an improved digital filter which no longer has zeros directly at $z = -1$. Since the low-pass filter has a gain of 1 at DC, we can set $G_0 = 1$. To find G_1, we must first find the gain of the original analog low-pass filter at the Nyquist frequency $f = F_s/2$ or $\omega = 2\pi f / F_s = \pi$, as demonstrated below:

$$G_1 = |H(\omega = \pi)| \tag{2.17}$$

$$G_1^2 = \frac{(\omega_N^2 - \omega_0^2)^2 + g^2 \omega_N^2 \omega_0^2 / Q^2}{(\omega_N^2 - \omega_0^2)^2 + \omega_N^2 \omega_0^2 / Q^2}, \quad \omega_N = \pi \tag{2.18}$$

Where Equation 2.18 shows the G_1 calculation for the parametric peaking filter. The other three parameters, A, B, and W, can be found by further constraining the magnitude response. Orfanidis uses the following constraints:

$$\frac{\delta}{\delta \omega} |H_d(\omega_0)|^2 = 0, \qquad |H_d(\omega_0)|^2 = g^2, \qquad |H_d(\omega_B)|^2 = G_B^2 \tag{2.19}$$

The first two equations constrain the peak of the filter. By taking the magnitude function's derivative and setting it to 0, we can ensure that the peak is at the correct location at $\omega = \omega_0$. To force the peak to also have the correct gain, we can set the magnitude response equal to the desired gain g at ω_0. The third equation controls the bandwidth by constraining the magnitude to be G_B at the desired frequency ω_B. G_B can be any arbitrary gain, so long as it is between unity and the peak gain g, for instance $G_B = \sqrt{g}$. Since the filter is symmetric around the cutoff frequency ω_0, there are two points on the magnitude response with gain G_B, ω_1 and ω_2. Orfanidis takes advantage of this fact to simplify the algebra required to solve for each parameter. When solving for these parameters, the constrained frequencies (ω_0, ω_1, and ω_2) can be pre-warped to ensure the peak and bandwidth are matched between the analog and digital domains.

After computing the five parameters and substituting them into Equation 2.14, the BLT can be taken to compute the digital filter coefficients. In Figure 2.13, when the cutoff frequency is half the Nyquist frequency, $f_0 = F_s/4$, the new filter's magnitude response is remarkably close to that of the analog filter's. In particular, notice how the gain near the Nyquist frequency has been corrected.

However, moving the cutoff frequency closer to Nyquist, say $f_0 = 15$kHz, we see odd behavior in Figure 2.14. Why does this happen? With a peak gain of 10dB, we can compute $G_B = \sqrt{g}$. Taking the square root in the linear domain is equivalent to dividing by two in the logarithmic domain, so $G_B = 5$dB. As discussed earlier, the magnitude response contains two

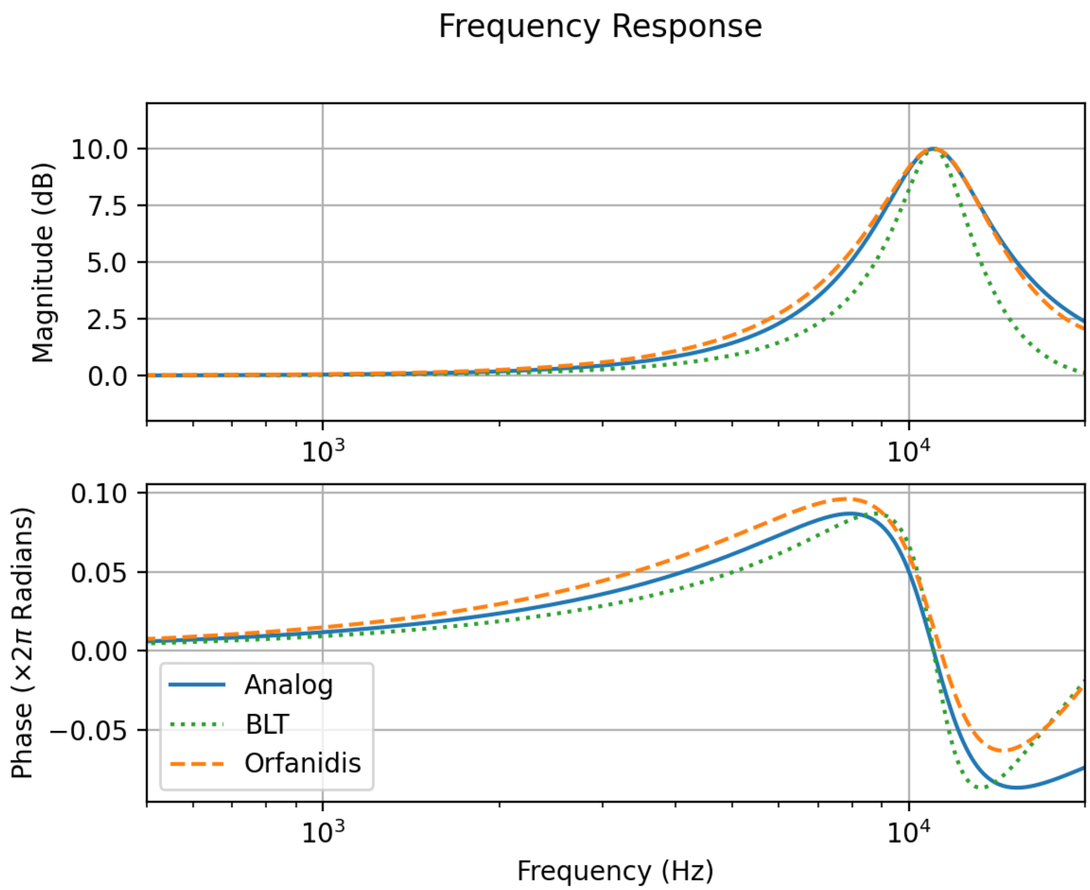

Figure 2.13. Peaking filter parameters: $F_s = 44.1\text{kHz}$, $f_0 = F_s/4$, $Q = 1.5$, $g = 10\text{dB}$

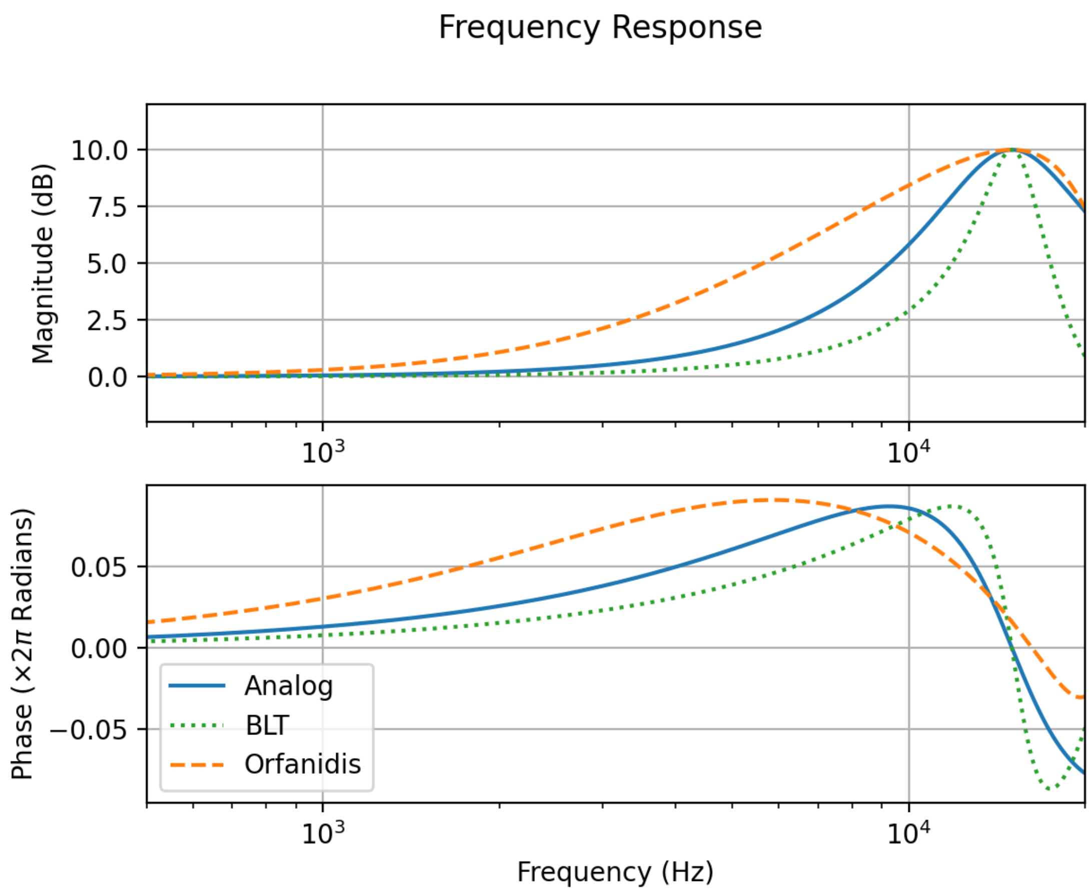

Figure 2.14. Peaking filter parameters: $F_s = 44.1\text{kHz}$, $f_0 = 15\text{kHz}$, $Q = 1$, $g = 10\text{dB}$

frequencies where the gain equals G_B, ω_1 and ω_2. Although not visible, the upper position ω_2 is above the Nyquist frequency. This becomes a problem, as attempting to pre-warp ω_2 results in undefined behavior and incorrect bandwidth computation. This effect becomes worse with smaller Q values and larger cutoff frequencies f_0.

While this approach maintains the filter order and a minimum-phase response, the troublesome bandwidth computation makes this filter difficult to use. In addition, this approach is specific to the peaking filter, as the constraints used do not work across the rest of the parametric filter family. Despite these shortcomings, this idea of prescribed Nyquist-frequency gain is worth further exploring and will be utilized in our proposed solution.

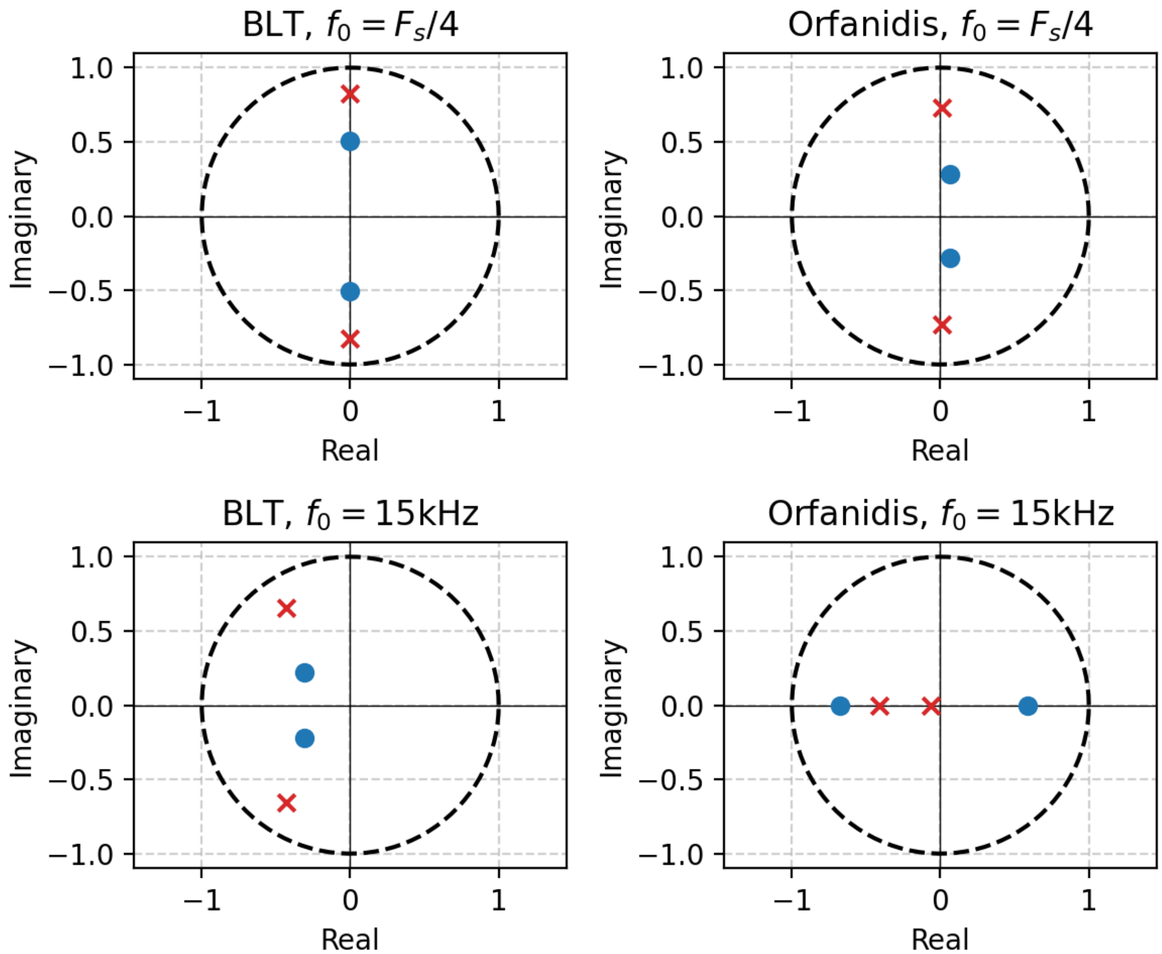

Figure 2.15. Pole-zero diagrams ($F_s = 44.1$kHz, $g = 10$dB).

2.5 Additional Methods

The works analyzed in this chapter do not form an exhaustive list of all strategies available for designing digital parametric filters. For instance, Harris and Booking's paper [5] details a design for deriving these filters directly in the digital domain; no need for the BLT or any other mapping. By exploiting destructive interference, this technique cleverly utilizes digital all-pass filters to achieve the desired magnitude responses. For instance, running an all-pass filter in parallel with the original signal and summing the two paths forms the basis for the peaking filter design. Unfortunately, the digital all-pass filters used suffer from a similar frequency warping issue, which extends to the derived parametric filters.

Another approach worth mentioning is that of Sarkka and Huovilainen [10]. Out of all the techniques discussed thus far, this paper derives digital filters which most closely match the frequency responses of their analog counterparts. Despite this, the dramatically increased filter order and complexity of computing the coefficients make it unideal for use in real-time audio processing applications.

Chapter 3

Proposed Solution

For the proposed solution, we will be building upon Orfanidis' work as discussed in Section 2.4. As a reminder, the key techniques employed are prescribed Nyquist-frequency gain and pre-warped frequency response constraints. However, we will be making three important alterations:

1. We will utilize a slightly different general second-order transfer function.
2. Improve the bandwidth constraint.
3. Apply variations of this approach to all parametric filters.

Through trial-and-error, we have found it to be easier to work with the following general transfer function:

$$H_d(s) = \frac{G_1 s^2 + sB\frac{\omega_d}{Q_d} + G_0\omega_d^2}{s^2 + s\frac{\omega_d}{Q_d} + \omega_d^2} \tag{3.1}$$

$$|H_d(\omega)|^2 = \frac{(G_1\omega^2 - G_0\omega_d^2)^2 + B^2\omega^2\omega_d^2/Q_d^2}{(\omega^2 - \omega_d^2)^2 + \omega^2\omega_d^2/Q_d^2} \tag{3.2}$$

We maintain the G_0, G_1, B, and $W \leftarrow \omega_d$ parameters as found in Orfanidis' method while introducing a new Q_d parameter. The generality of the transfer function is maintained, as each coefficient can still be uniquely set, while bringing it more in line with the peaking filter transfer function. This change allows for more intuitive reasoning about how each parameter affects the frequency response.

As explained in Section 2.4, the G_0 and G_1 parameters correspond to the magnitude of the frequency response at DC and Nyquist frequency, respectively. Q_d can be thought of as the corrected quality factor, B as the adjusted gain factor, and ω_d as the corrected cutoff frequency, similar to the parameters of Sierra's method (Section 2.3). When computing these parameters, it is important to use pre-warped frequency values, as discussed in Section 2.4. This will ensure that after applying the BLT, the filter characteristics are preserved.

Taking the bilinear transform of the general transfer function in Equation 3.1 generates the following digital filter coefficients. Equation 2.6 can be used to derive the time-domain difference equation representation.

$$\begin{aligned} b_0 &= n^2 G_1 + nB/Q_d + G_0, \quad & b_1 &= 2(G_0 - n^2 G_1), \quad & b_2 &= n^2 G_1 - nB/Q_d + G_0, \\ a_0 &= n^2 + n/Q_d + 1, & a_1 &= 2(1 - n^2), & a_2 &= n^2 - n/Q_d + 1, \end{aligned}$$

$$\text{where } n = \frac{\omega_0}{\omega_d} \cdot \frac{1}{\tan(\omega_0/2)}$$

3.1 Peaking Filter

Since all the existing techniques focus on the peaking filter, we will begin here. Instead of using the peaking transfer function from the Audio EQ Cookbook [2], we will use the slightly modified version from Section 2.1:

$$H(s) = \frac{s^2 + sg\frac{\omega_0}{Q} + \omega_0^2}{s^2 + s\frac{\omega_0}{Q} + \omega_0^2} \tag{3.3}$$

$$|H(\omega)|^2 = \frac{(\omega^2 - \omega_0^2)^2 + g^2\omega^2\omega_0^2/Q^2}{(\omega^2 - \omega_0^2)^2 + \omega^2\omega_0^2/Q^2} \tag{3.4}$$

To obtain G_0 and G_1, we must find the gain of $H(s)$ at DC ($\omega = 0$) and Nyquist ($\omega = \pi$). Note that we do not pre-warp the frequencies here.

$$G_0 = |H(\omega = 0)| = 1 \tag{3.5}$$

$$G_1 = |H(\omega = \pi)| = \sqrt{\frac{(\omega_N^2 - \omega_0^2)^2 + g^2\omega_N^2\omega_0^2/Q^2}{(\omega_N^2 - \omega_0^2)^2 + \omega_N^2\omega_0^2/Q^2}}, \quad \omega_N = \pi \tag{3.6}$$

A key characteristic of the peaking filter is, of course, its peak. We know that this peak is located at ω_0 with a gain of g. We can use this as a constraint to solve for B:

$$|H_d(\omega = \omega_0)|^2 = g^2 \tag{3.7}$$

$$\Rightarrow B^2 = \frac{\omega_0^2\omega_d^2 g^2 + Q_d^2 g^2(\omega_0^2 - \omega_d^2)^2 - Q_d^2(G_0\omega_d^2 - G_1\omega_0^2)^2}{\omega_0^2\omega_d^2} \tag{3.8}$$

Although this equation constrains the gain to be correct at cutoff frequency ω_0, it does not guarantee that the peak is at ω_0. To center the peak, we can set the derivative of the magnitude response to be 0 at this location, ensuring an extrema. By substituting B with Equation 3.8 and $\omega \leftarrow \omega_0$, we can solve for ω_d:

$$\frac{\delta}{\delta\omega}|H(\omega)|^2\bigg|_{\omega=\omega_0,B} = 0 \tag{3.9}$$

$$\Rightarrow \omega_d^2 = \omega_0^2\sqrt{\frac{g^2 - G_1^2}{g^2 - G_0^2}} \tag{3.10}$$

Lastly, to solve for Q_d, we can apply a constraint to the bandwidth. Similar to Orfanidis, we can choose an arbitrary gain between unity and peak gain g. In this case, we will use $G_B = \sqrt{g}$. As discussed in Section 2.4, the peaking filter contains two distinct points with gain G_B. As a reminder, Orfanidis uses both these points to formulate a constraint, which we demonstrated results in strange behavior in the frequency response. To prevent this, we will only utilize the lower of the two points; this lower point is guaranteed to always be below the cutoff frequency, which will remain below the Nyquist limit. To utilize G_B, we must first find its location ω_B in the analog filter $H(s)$.

$$|H(\omega = \omega_B)|^2 = g \tag{3.11}$$

$$\Rightarrow \omega_B^2 = \omega_0^2 \cdot \frac{2Q^2 + g - \sqrt{g(4Q^2+g)}}{2Q^2} \tag{3.12}$$

Before moving onto the next step, ω_B must be pre-warped to preserve the bandwidth in anticipation of the BLT. We can solve for Q_d by constraining the gain to be G_B at the pre-warped ω_B.

$$|H_d(\omega = \omega_B)|^2 = g \tag{3.13}$$

$$\Rightarrow Q_d^2 = \omega_B^2 \omega_d^2 \cdot \frac{g - B^2}{G_0^2\omega_d^4 - 2G_0G_1\omega_B^2\omega_d^2 + G_1^2\omega_B^4 - \omega_B^4 g + 2\omega_B^2\omega_d^2 g - \omega_d^4 g} \tag{3.14}$$

To remove the dependency on B, we can use Equation 3.8 and solve for Q_d.

$$\begin{aligned}\Rightarrow Q_d^2 = \omega_B^2\omega_0^2\omega_d^2 g(1-g)/(&G_0^2\omega_0^2\omega_d^4 - 2G_0G_1\omega_B^2\omega_0^2\omega_d^2 + G_1^2\omega_B^4\omega_0^2 - \omega_B^4\omega_0^2 g + \\ &+ 2\omega_B^2\omega_0^2\omega_d^2 g + \omega_B^2 g^2(\omega_0^2 - \omega_d^2)^2 - \omega_B^2(G_0\omega_d^2 - G_1\omega_0^2)^2 - \omega_0^2\omega_d^4 g)\end{aligned} \tag{3.15}$$

Note that similar to ω_B, ω_0 and ω_d must be pre-warped. With that, each parameter has been accounted for. Since the computation of certain parameters depends on others, it is important to process them in the correct order:

1. Calculate G_0 and G_1 using the original analog parametric magnitude response $H(s)$ as shown in Equation 3.5 and Equation 3.6.
2. Compute ω_d using Equation 3.10.
3. Find ω_B using Equation 3.12.
4. Compute Q_d using Equation 3.15.
5. Compute B using Equation 3.8.
6. Apply the BLT to derive the digital filter coefficients using the table found at the beginning of this chapter.

Let's examine the newly derived filter. Setting the cutoff frequency to be near the Nyquist limit, the frequency response plot in Figure 3.1 shows that the magnitude response is remarkably close to that of the analog filter and a large improvement over the standard BLT-derived filter. Not only is this achieved without increasing the order of the filter, but the pole-zero diagram in Figure 3.2 demonstrates that the filter remains minimum-phase, stable, and causal.

Worth noting is that as the cutoff frequency moves away from Nyquist and the bandwidth becomes more narrow, the proposed filter approaches the standard BLT-derived filter, as seen in Figure 3.3. In this cases, G_1 approaches 1, resulting in the transfer function reverting back to the original peaking filter transfer function. This behavior is desired, as under these conditions the standard BLT-derived frequency response approaches the original analog response.

In the opposite extreme, as the cutoff frequency approaches Nyquist and the bandwidth becomes wider, the proposed and standard BLT filters diverge, as seen in Figure 3.5. The proposed solution maintains a nearly matching magnitude response while the standard BLT-derived filter diverges dramatically.

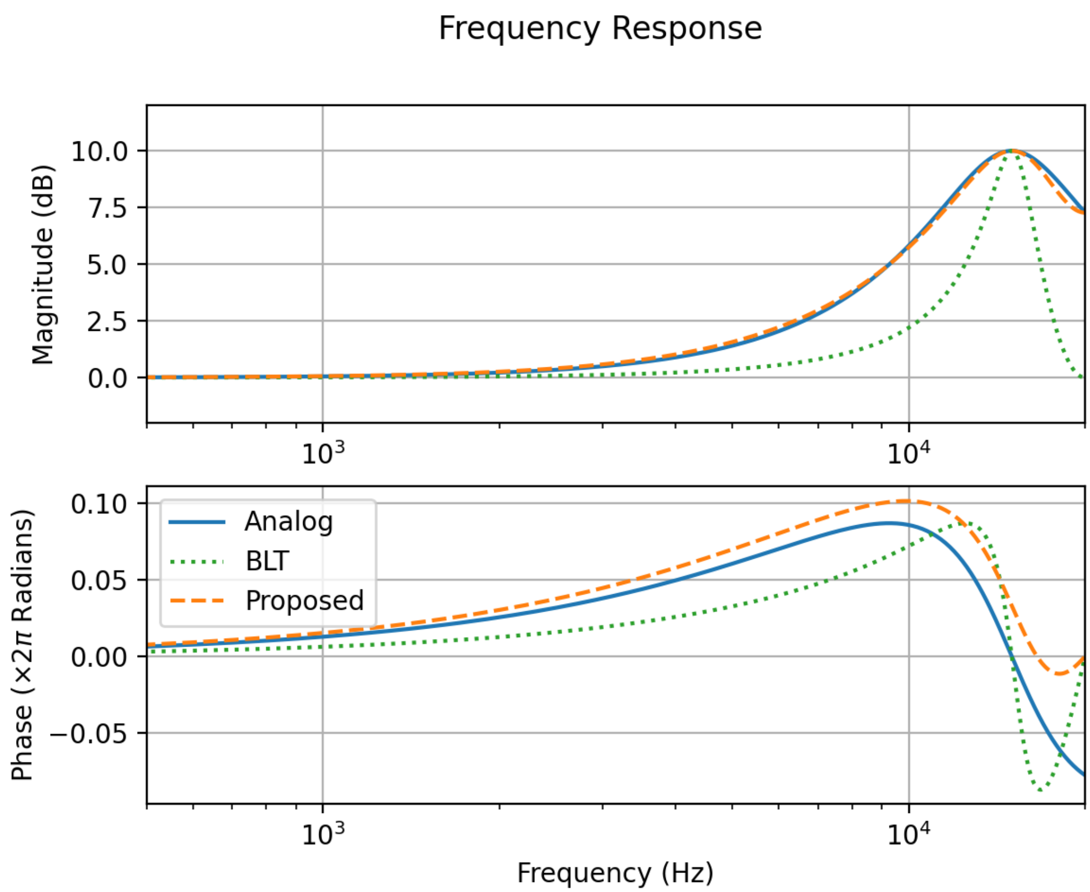

Figure 3.1. Proposed method to derive the peaking filter ($F_s = 40\text{kHz}$, $f_0 = 15\text{kHz}$, $Q = 1$, $g = 10\text{dB}$).

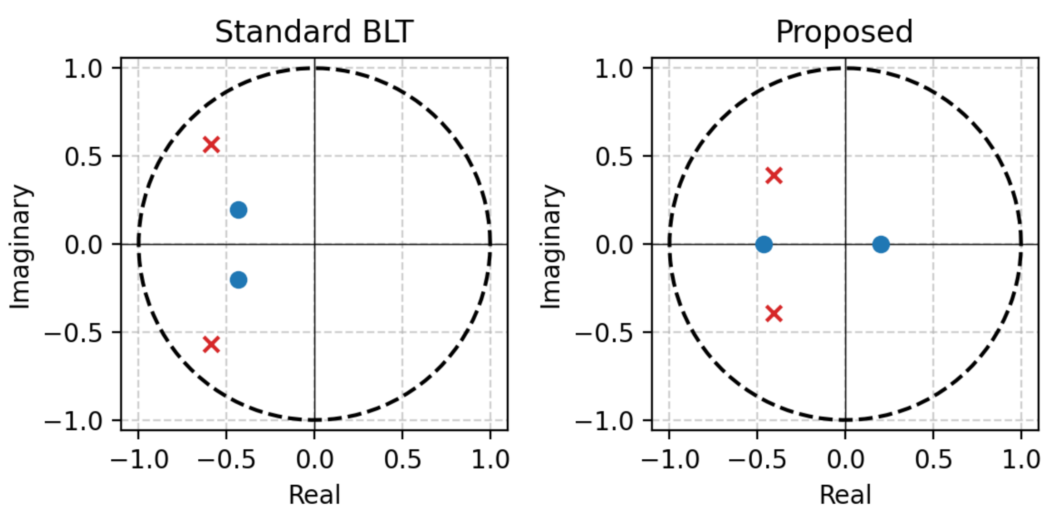

Figure 3.2. Pole-zero diagrams of peaking filters ($F_s = 40\text{kHz}$, $f_0 = 15\text{kHz}$, $Q = 1$, $g = 10\text{dB}$).

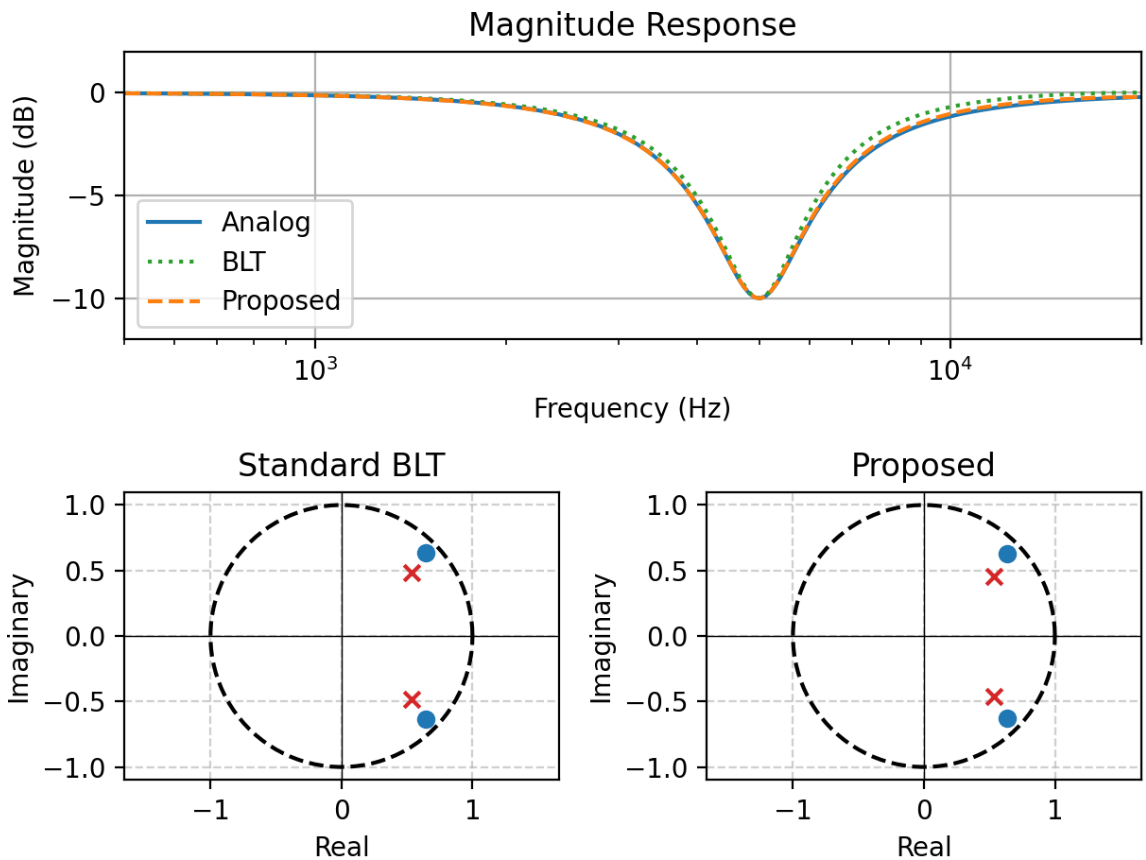

Figure 3.3. Peaking filter with low cutoff frequency and narrow bandwidth ($F_s = 40$kHz, $f_0 = 5$kHz, $Q = 2$, $g = -10$dB).

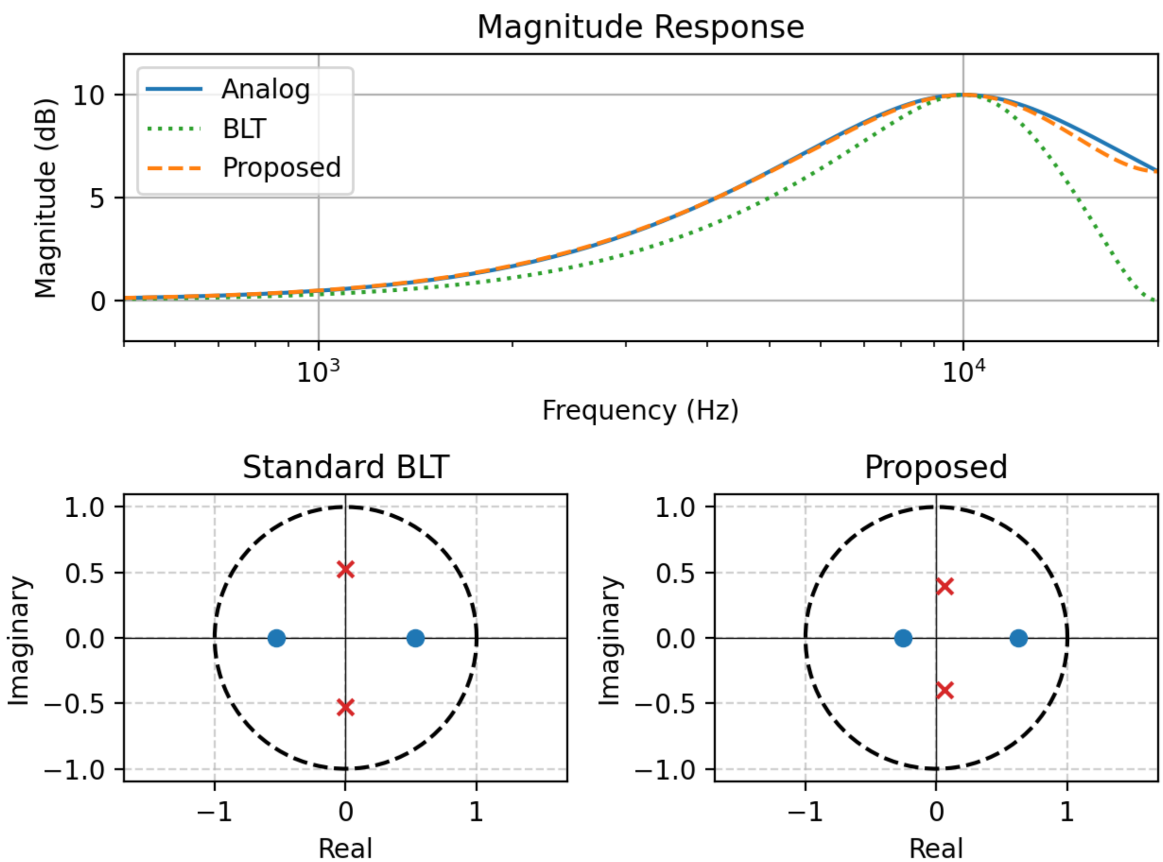

Figure 3.4. Peaking filter with cutoff at half Nyquist and medium bandwidth ($F_s = 40$kHz, $f_0 = 10$kHz, $Q = 1$, $g = 10$dB).

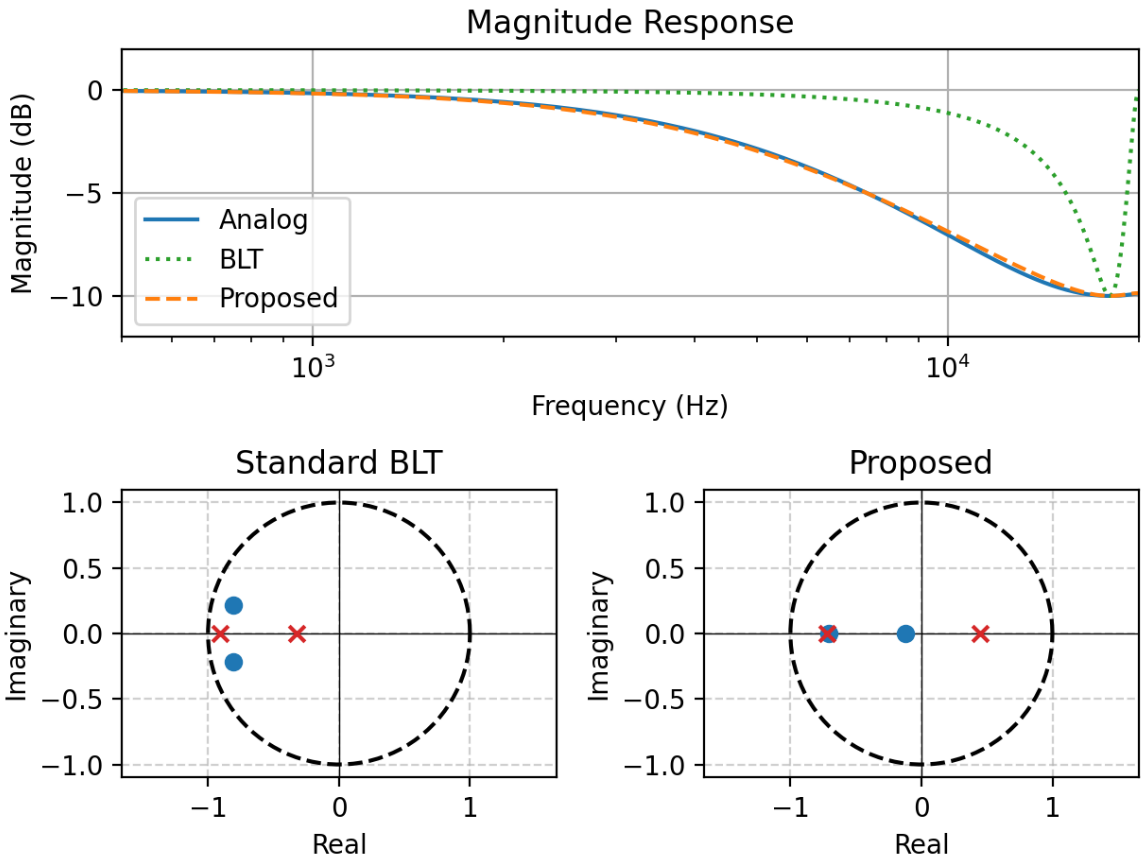

Figure 3.5. Peaking filter with very high cutoff frequency and wide bandwidth ($F_s = 40$kHz, $f_0 = 18$kHz, $Q = 0.5$, $g = -10$dB).

3.2 Band-pass Filter

The analog parametric band-pass filter has the following transfer function [2]:

$$H(s) = \frac{s\frac{\omega_0}{Q}}{s^2 + s\frac{\omega_0}{Q} + \omega_0^2} \tag{3.16}$$

$$|H(\omega)|^2 = \frac{\omega^2\omega_0^2/Q^2}{(\omega^2-\omega_0^2)^2 + \omega^2\omega_0^2/Q^2} \tag{3.17}$$

There are many similarities between this band-pass filter and the peaking filter from Section 3.1. In particular, the peak/maxima of the filter is located at ω_0 but always has a gain of 0dB. In addition, the Q parameter adjusts the bandwidth. Due to these similarities, we can apply the same process to derive the digital filter. We begin with G_0 and G_1.

$$G_0 = |H(\omega=0)| = 0 \tag{3.18}$$

$$G_1 = |H(\omega=\pi)| = \sqrt{\frac{\omega_N^2\omega_0^2/Q^2}{(\omega_N^2-\omega_0^2)^2 + \omega_N^2\omega_0^2/Q^2}}, \quad \omega_N = \pi \tag{3.19}$$

Instead of deriving B from scratch, we can borrow the B derivation from peaking filter Equation 3.8. Since the band-pass filter has a peak gain 0dB, we use $g = 1$.

$$B^2 = 1 - G_1^2 Q_d^2 \frac{\omega_0^2}{\omega_d^2} + Q_d^2 \frac{(\omega_0^2-\omega_d^2)^2}{\omega_0^2\omega_d^2} \tag{3.20}$$

Similarly, we can use Equation 3.10 to find ω_d. We can further simplify this expression by using $g = 1$ and $G_0 = 0$.

$$\omega_d^2 = \omega_0^2\sqrt{1-G_1^2} \tag{3.21}$$

Lastly, we use Equation 3.15 to obtain Q_d utilizing the previously used substitutions to simplify. To compute ω_B, we find the frequency at which the magnitude $|H(\omega)|$ is equal to arbitrary bandwidth gain G_B. Unlike in the peaking filter, we do not use $G_B = \sqrt{g}$, as $g = 1$. Through trial and error, we have found $G_B = \frac{1}{2}$ to provide good results.

$$\omega_B^2 = \omega_0^2 \cdot \frac{1 + G_B^2(2Q^2 - 1) - \sqrt{-(G_B^2 - 1)(4Q^2G_B^2 - G_B^2 + 1)}}{2Q^2G_B^2} \tag{3.22}$$

$$Q_d^2 = \omega_B^2\omega_0^2\omega_d^2(G_B^2 - 1)/(G_1^2\omega_B^4\omega_0^2 - G_1^2\omega_B^2\omega_0^4 - \omega_B^4\omega_0^2G_B^2 + \omega_B^2\omega_0^4 + 2\omega_B^2\omega_0^2\omega_d^2G_B^2 - \\ - 2\omega_B^2\omega_0^2\omega_d^2 + \omega_B^2\omega_d^4 - \omega_0^2\omega_d^4G_B^2) \tag{3.23}$$

Following the steps outlined in Section 3.1, we can compute the digital band-pass filter coefficients. Comparing these results to that of the standard BLT-derived filter in Figure 3.6, the proposed approach is able to more closely match the magnitude response of the original analog filter. As with the peaking filter, the pole-zero diagram in Figure 3.7 demonstrates that the band-pass filter remains minimum-phase, causal, and stable.

Notice that in the pole-zero diagram, the standard BLT-derived filter contains a zero at $z = 1$ and a second at $z = -1$, implying the original analog filter had zeros located at $s = 0$ and $s \to \infty$. This is logical, as a band-pass filter should attenuate frequencies distant from ω_0. However, due to there being a zero at $z = -1$ which corresponds to the Nyquist frequency, the gain at Nyquist will always be 0, which deviates from its analog counterpart. Compare this to the pole-zero diagram of our proposed filter, which maintains the zero at DC, $z = 1$, but moves the second zero off the unit circle, preventing the magnitude response from reaching 0 at Nyquist. The poles are also adjusted in response to the zero movement to maintain the bandwidth and center frequency ω_0.

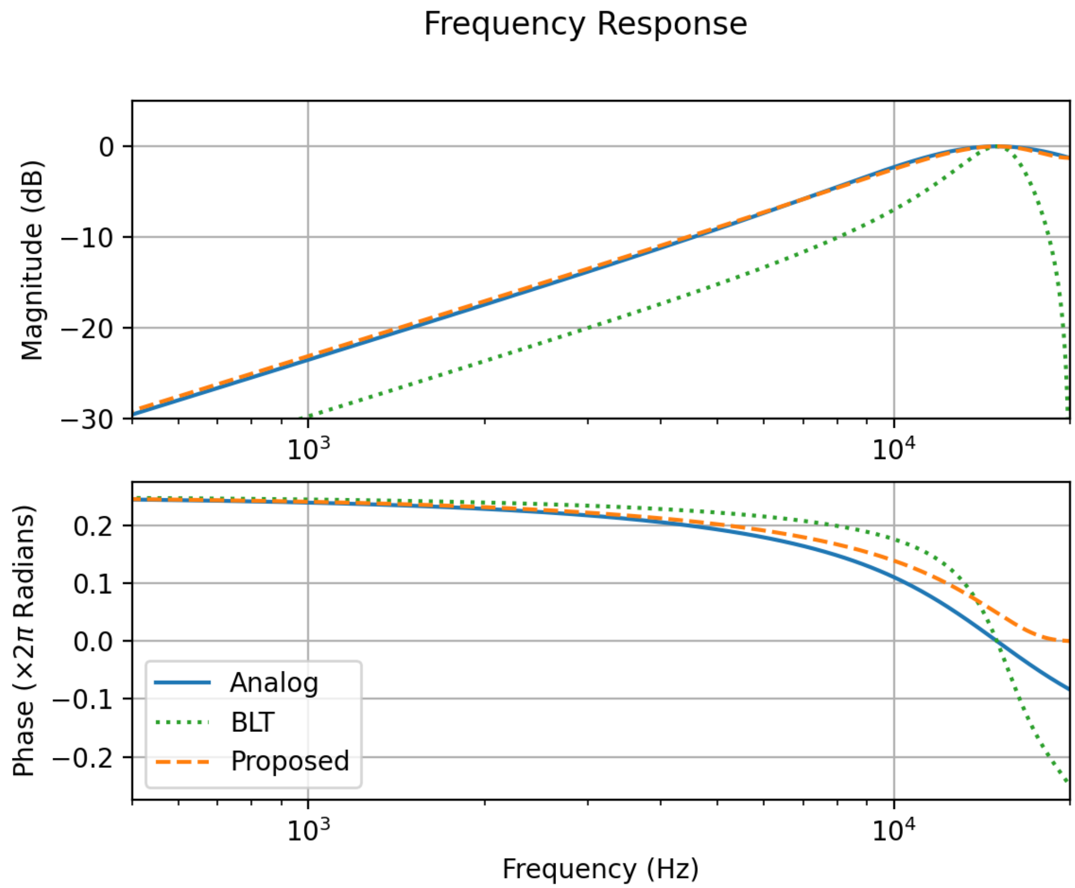

Figure 3.6. Proposed method to derive band-pass filter ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 1$).

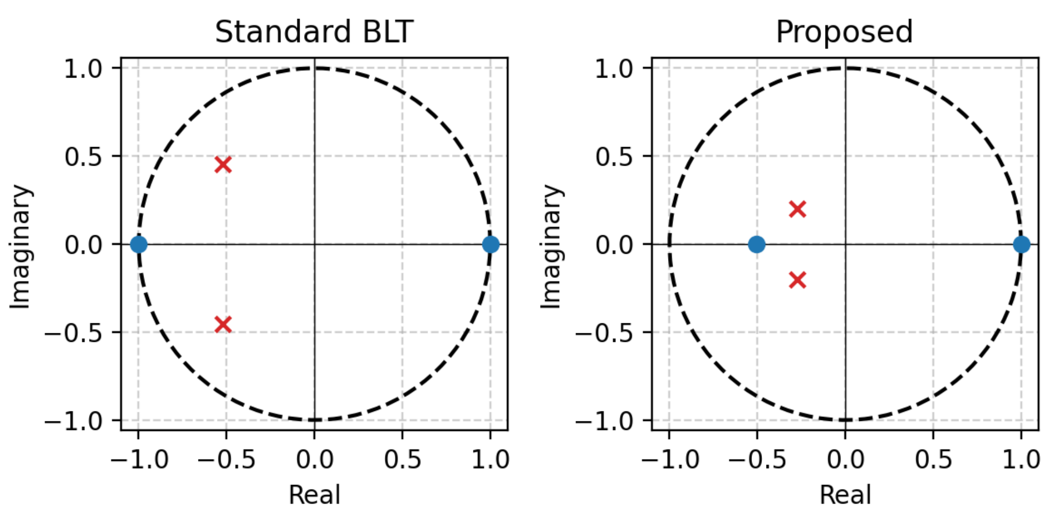

Figure 3.7. Pole-zero diagrams of band-pass filters ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 1$).

3.3 Notch Filter

Like the band-pass filter, we can take advantage of the notch filter's similarities to the peaking filter. Namely, the peak/extrema is located at cutoff frequency ω_0, but with a gain of $g = 0$. The quality Q continues to control the filter bandwidth. The parametric notch filter transfer function is shown below [2].

$$H(s) = \frac{s^2 + \omega_0^2}{s^2 + s\frac{\omega_0}{Q} + \omega_0^2} \tag{3.24}$$

$$|H(\omega)|^2 = \frac{(\omega^2 - \omega_0^2)^2}{(\omega^2 - \omega_0^2)^2 + \omega^2\omega_0^2/Q^2} \tag{3.25}$$

Starting with G_0 and G_1, these can be computed using $|H(s)|$ in Equation 3.25.

$$G_0 = |H(\omega = 0)| = 1 \tag{3.26}$$

$$G_1 = |H(\omega = \pi)| = \sqrt{\frac{(\omega_N^2 - \omega_0^2)^2}{(\omega_N^2 - \omega_0^2)^2 + \omega_N^2\omega_0^2/Q^2}}, \quad \omega_N = \pi \tag{3.27}$$

For the peaking and band-pass filters, we purposely adjusted any zeros at Nyquist $z = -1$ by moving them off the unit circle. In this case, there are no zeros at Nyquist. In fact, both zeros are located at the cutoff frequency ω_0. This is desired behavior and changes to the zeros would cause the notch location to change. To maintain the zero positions, we must set $B = 0$, as B corresponds to the extrema gain. To compute ω_d and Q_d we can reuse the calculations from Section 3.1. Additional simplifications can be applied by using $g = 0$ and $G_0 = 1$.

$$\omega_d^2 = \omega_0^2 G_1 \tag{3.28}$$

$$\omega_B^2 = \omega_0^2 \cdot \frac{Q^2G_B^2 - Q^2 - \frac{1}{2}G_B^2 + \frac{1}{2}G_B\sqrt{-4Q^2G_B^2 + 4Q^2 + G_B^2}}{Q^2(G_B^2 - 1)} \tag{3.29}$$

$$Q_d^2 = \frac{\omega_B^2\omega_d^2G_B^2}{(G_1\omega_B^2 - \omega_d^2)^2 - G_B^2(\omega_B^2 - \omega_d^2)^2} \tag{3.30}$$

Once again, through trial and error, we have found $G_B = \frac{1}{2}$ to provide good results. Following the process outlined in Section 3.1, the digital filter coefficients can be computed. The resulting magnitude response is a significant improvement on the standard BLT-derived filter, closely matching the analog response as seen in Figure 3.8. The pole-zero diagram in Figure 3.9 also demonstrates that the proposed solution remains minimum-phase, causal, and stable. As predicted, the zeros remain at the same positions on the unit circles between the two digital filters; the poles are entirely responsible for the improved frequency response.

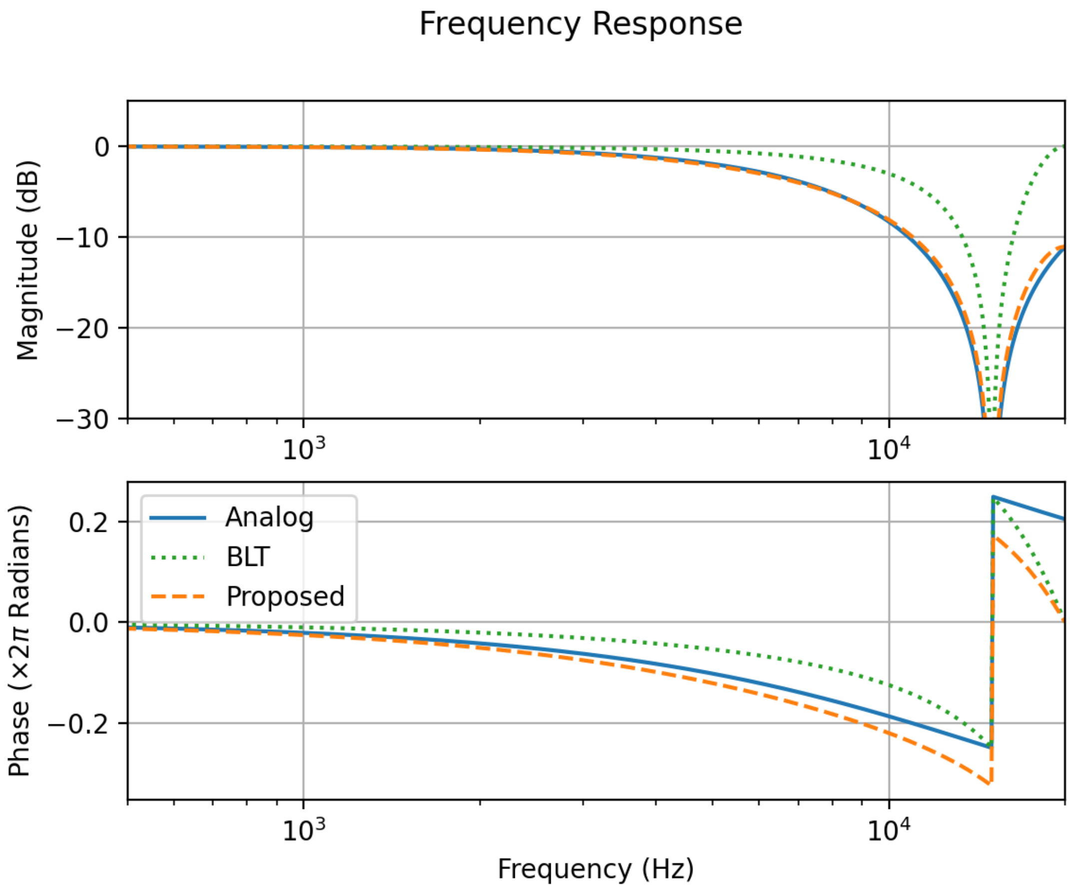

Figure 3.8. Proposed method to derive the notch filter ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 0.5$).

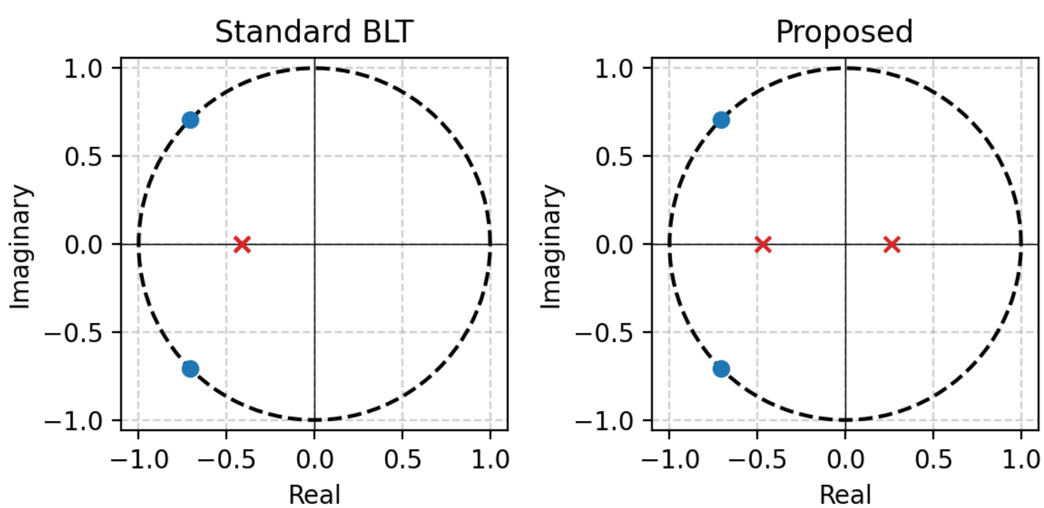

Figure 3.9. Pole-zero diagrams of notch filters ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 0.5$).

3.4 High-pass Filter

Although the transfer function of the high-pass filter shown below [2] may appear similar to that of the peaking filter, it shares a key limitation with the notch filter: the rigid placement of its zeros reduces the number of available degrees of freedom. Specifically, the high-pass filter requires both of its zeros to be fixed at DC ($s = 0$ or $z = 1$). Even slight adjustments to the zero placement affects the slope of the frequency response. As a result, the filter must satisfy $G_0 = 0$ and $B = 0$, which constrains the design parameters to just G_1, Q_d, and ω_d.

$$H(s) = \frac{s^2}{s^2 + s\frac{\omega_0}{Q} + \omega_0^2} \tag{3.31}$$

$$|H(\omega)|^2 = \frac{\omega^4}{(\omega^2 - \omega_0^2)^2 + \omega^2\omega_0^2/Q^2} \tag{3.32}$$

As discussed, $G_0 = 0$. The value of G_1 can be determined using the method described in Section 3.1.

$$G_1 = |H(\omega = \pi)| = \omega_N^2 \cdot \sqrt{\frac{1}{(\omega_N^2 - \omega_0^2)^2 + \omega_N^2\omega_0^2/Q^2}}, \quad \omega_N = \pi \tag{3.33}$$

To derive ω_d we can avoid the need to solve a system of equations by instead examining the slope of the filter. Even for high cutoff frequencies, the majority of the sloped region will be within the approximately linear mapping of the BLT. As a result, to preserve the filter slope, we must avoid pre-warping the cutoff frequency. In addition, to account for the adjusted filter gain via G_1, the cutoff must be multiplied by a factor of G_1, resulting in the following formula for ω_d.

$$\omega_d = \omega_0\sqrt{G_1} \tag{3.34}$$

When attempting to compute Q_d, we run into a problem. Unlike the previously derived filters, the high-pass filter does not always contain a peak/extrema. A Butterworth Filter can be obtained by setting $Q = 1/\sqrt{2}$. So a peak only exists when $Q > 1/\sqrt{2}$. To work around this issue, we can solve for Q_d in two different ways: once when $Q > 1/\sqrt{2}$ and again when $Q \leq 1/\sqrt{2}$.

3.4.1 Solving with a peak

Beginning with the simpler case of $Q > 1/\sqrt{2}$, we can utilize the previously used strategies to solve for Q_d. In particular, our goal is to match the peak of the digital high-pass filter to its analog counterpart. Due to our inability to adjust the zeros, we lack the degrees of freedom necessary to match both the peak location and gain. So, we have elected to match only the peak gain.

Unlike the previously discussed filters, the high-pass peak is not located at ω_0. To find the peak frequency, we can set the derivative of the magnitude response to 0 and solve for ω. We will call this frequency ω_p.

$$\frac{\delta}{\delta \omega_p} |H(\omega_p)|^2 = 0 \tag{3.35}$$

$$\Rightarrow \omega_p^2 = \omega_0^2 \cdot \frac{2Q^2}{2Q^2 - 1} \tag{3.36}$$

Next, we can substitute $\omega \leftarrow \omega_p$ back into the magnitude response to find G_p, the gain of the peak.

$$G_p^2 = |H(\omega)|^2 \big|_{\omega=\omega_p} = \frac{4Q^4}{4Q^2 - 1} \tag{3.37}$$

To be able to control the peak gain of the digital filter, we must apply the same process to $H_d(s)$.

$$\frac{\delta}{\delta \omega_{pd}} |H_d(\omega_{pd})|^2 = 0 \tag{3.38}$$

$$\Rightarrow \omega_{pd}^2 = \omega_d^2 \cdot \frac{2Q_d^2}{2Q_d^2 - 1} \tag{3.39}$$

$$G_{pd}^2 = |H_d(\omega)|^2\big|_{\omega=\omega_p} = \frac{4G_1^2 Q_d^4}{4Q_d^2 - 1} \tag{3.40}$$

Lastly, to match the peaks, we can set $G_p^2 = G_{pd}^2$ and solve for Q_d.

$$Q_d^2 = Q^2 \cdot \frac{2Q^2 + \sqrt{G_1^2 + 4Q^4 - 4G_1^2 Q^2}}{G_1^2(4Q^2 - 1)} = Q^2 \cdot \frac{G_p^2}{G_1^2}, \tag{3.41}$$

$$\text{where } G_p^2 = \frac{2Q^2 + \sqrt{G_1^2 + 4Q^4 - 4G_1^2 Q^2}}{4Q^2 - 1}$$

3.4.2 Solving with no peak

To solve for Q_d when $Q \leq 1/\sqrt{2}$, we have to employ a technique that does not utilize peak gain. After experimenting with different approaches, simply matching the gain at the digital cutoff frequency ω_d seems to provide good results. Since ω_d has not been pre-warped, we have to use the inverse of the pre-warping equation to find the equivalent frequency in the analog domain ω_a. By substituting $\omega \leftarrow \omega_a$ in $|H(\omega)|$, we can find gain G_d at ω_d.

$$\omega_a = 2\arctan(\omega_d/2) \tag{3.42}$$

$$G_d^2 = |H(\omega)|^2\big|_{\omega=\omega_a} = \frac{\omega_0^4}{(\omega_a^2 - \omega_0^2)^2 + \omega_a^2\omega_0^2/Q^2} \tag{3.43}$$

Similarly, we can find the gain of the digital high-pass filter $H_d(s)$ at frequency ω_d by $\omega \leftarrow \omega_d$ in $|H_d(\omega)|$. Solving for Q_d will allow us to adjust the gain.

$$G_d = |H_d(\omega)|\big|_{\omega=\omega_d} = Q_d G_1 \tag{3.44}$$

$$\Rightarrow Q_d = \frac{G_d}{G_1} \tag{3.45}$$

3.4.3 Combining the solutions

We can now compute Q_d for any value of Q. The solutions can be combined into the equation below, where G_p and G_d can be computed via Equations 3.41 and 3.43, respectively. At the $Q = 1/\sqrt{2}$ boundary, there exists a discontinuity. To mitigate this, we can linearly interpolate between the values. In the figures presented below, the interpolation occurs for values $1 < Q < 2$.

$$Q_d = \begin{cases} Q \cdot \frac{G_p}{G_1} & \text{if } Q > 1/\sqrt{2} \\ \frac{G_d}{G_1} & \text{if } Q \leq 1/\sqrt{2} \end{cases} \tag{3.46}$$

To obtain the digital filter coefficients, we once again take the BLT of Equation 3.1. The resulting filter has a magnitude response that closely matches the analog response, as seen in Figure 3.10. As predicted, the pole-zero diagram in Figure 3.11 displays the two zeros being located at $z = 1$ for both digital filters. Only the poles have been adjusted in the proposed solution. Figure 3.12 and Figure 3.13 demonstrate the proposed filter maintains the desired magnitude response for small and large values of Q. Note that while the peak gain is maintained, the peak frequency of the high Q digital response diverges from that of the analog response. Due to the constraint imposed on the zeros, it is impossible to also match the peak location without increasing the filter order. Instead, we compromise by favoring a correct filter slope and peak gain over peak location. Even with this trade-off, our proposed solution provide a significant improvement over the standard BLT-derived filter.

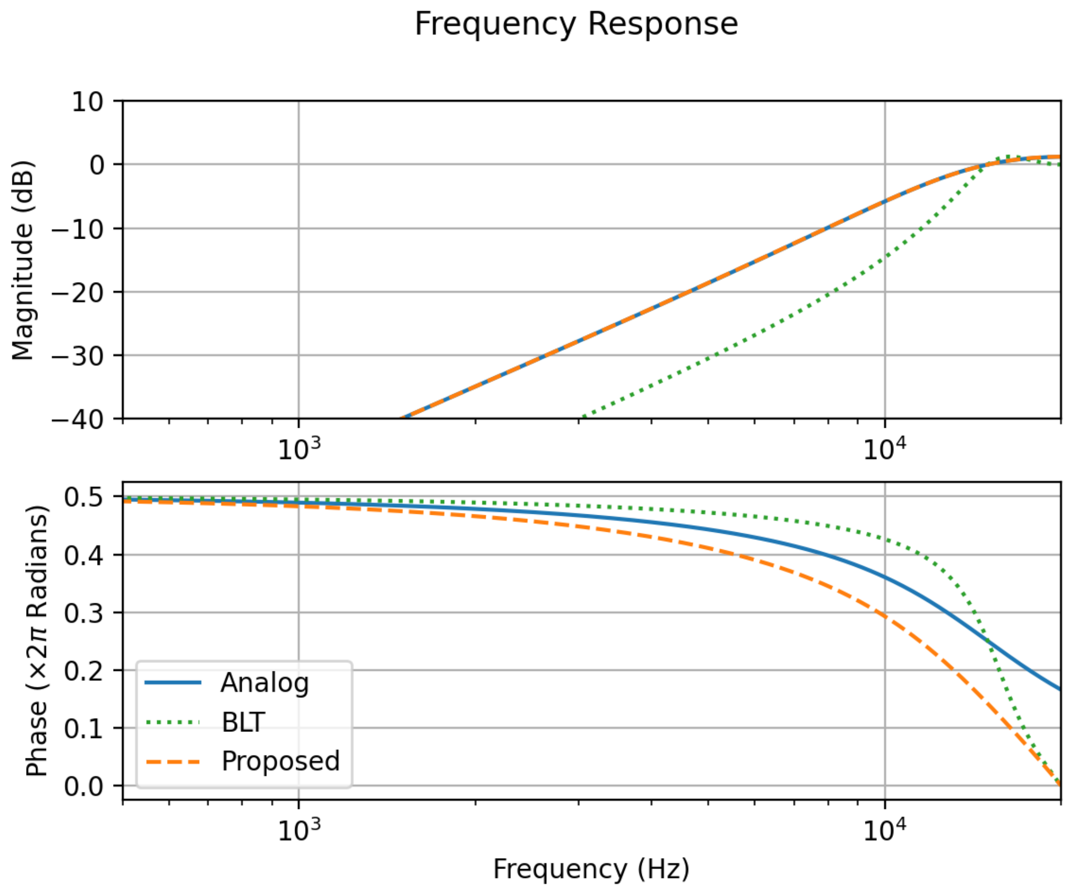

Figure 3.10. Proposed method to derive the high-pass filter ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 1$).

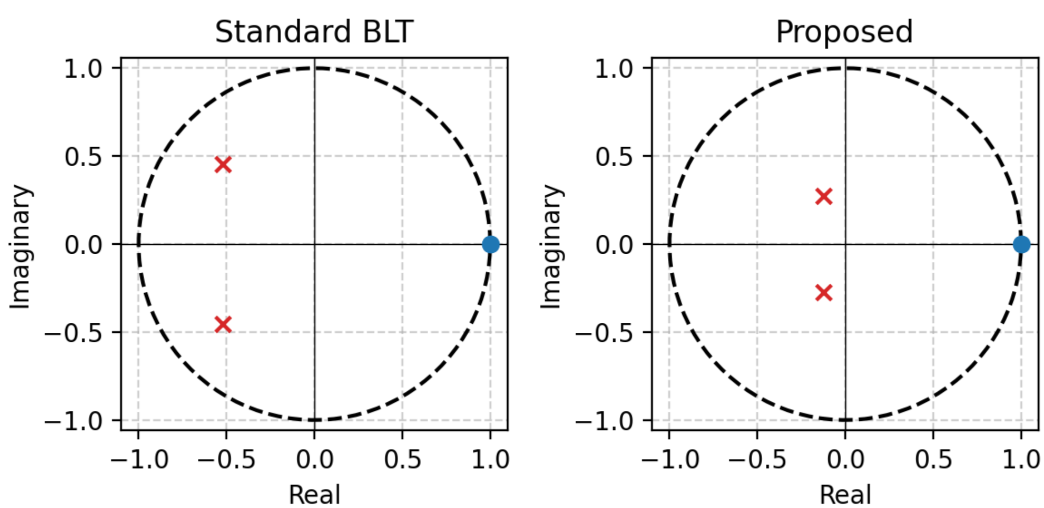

Figure 3.11. Pole-zero diagrams of high-pass filters ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 1$).

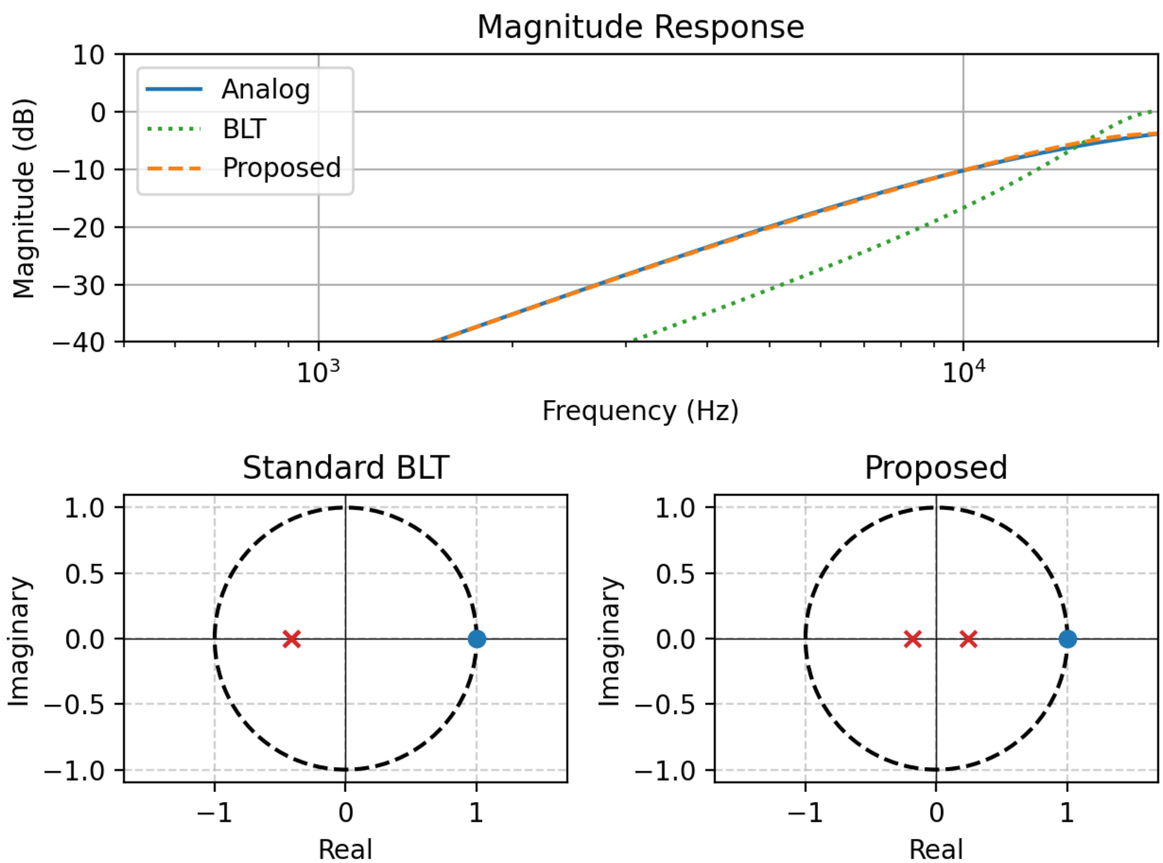

Figure 3.12. High-pass with small Q ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 0.5$).

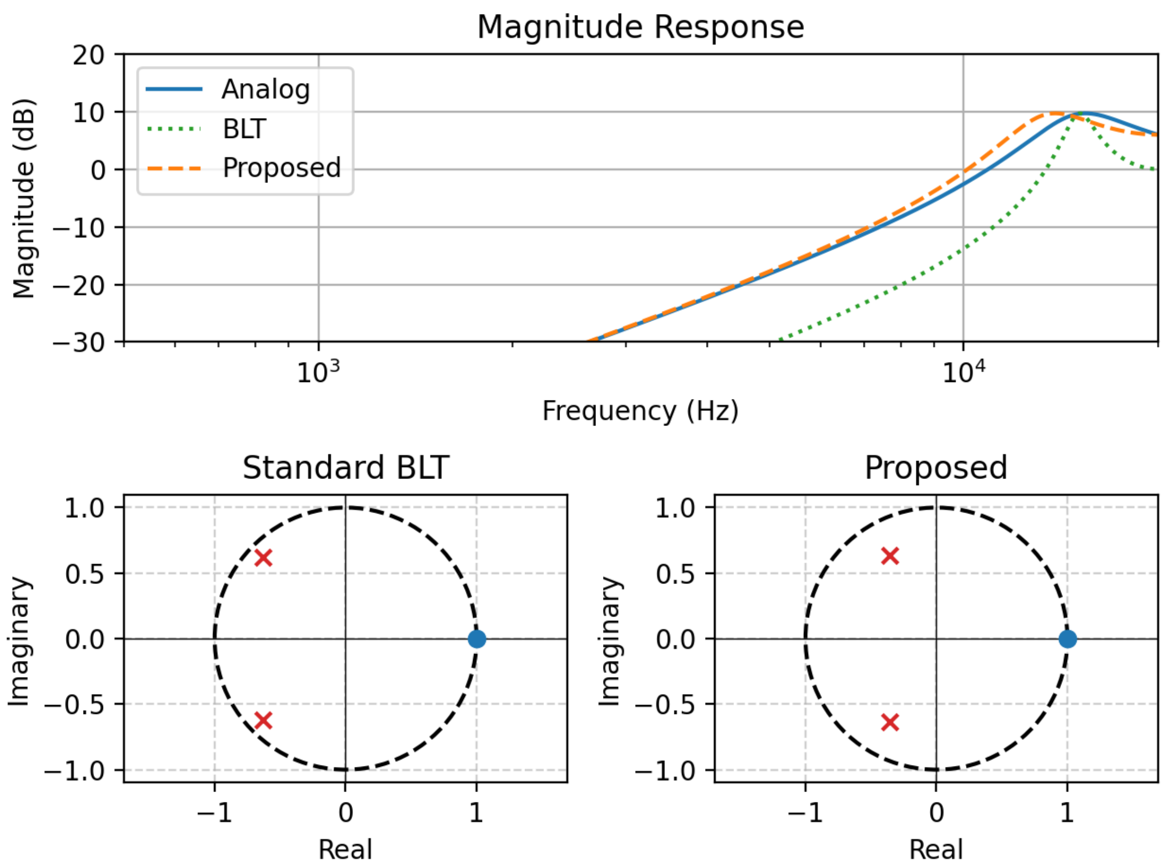

Figure 3.13. High-pass with large Q ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 3$).

3.5 Low-pass Filter

Similar to the high-pass filter, the parametric low-pass filter magnitude response only has a peak when $Q > 1/\sqrt{2}$, where $Q = 1/\sqrt{2}$ corresponds to the Butterworth low-pass filter. As a result, we must solve for the parameters in two scenarios, depending on whether the peak is present. Below is the analog parametric low-pass filter [2]:

$$H(s) = \frac{\omega_0^2}{s^2 + s\frac{\omega_0}{Q} + \omega_0^2} \tag{3.47}$$

$$|H(\omega)|^2 = \frac{\omega_0^4}{(\omega^2 - \omega_0^2)^2 + \omega^2\omega_0^2/Q^2} \tag{3.48}$$

Once again, we can compute G_0 and G_1 by computing the magnitude at DC and the Nyquist frequency:

$$G_0 = |H(\omega = 0)| = 1 \tag{3.49}$$

$$G_1 = |H(\omega = \pi)| = \sqrt{\frac{\omega_0^4}{(\omega_N^2 - \omega_0^2)^2 + \omega_N^2\omega_0^2/Q^2}}, \quad \omega_N = \pi \tag{3.50}$$

3.5.1 Solving with a peak

Starting with the scenario where $Q > 1/\sqrt{2}$, a peak is present in the magnitude response. As a result, we can use a similar process for deriving the peaking filter in Section 3.1. In particular, we can solve for B and ω_d by constraining the peak frequency and gain. First, we need to find the peak location ω_p in the analog prototype by setting the magnitude derivative equal to zero. The analog peak frequency ω_p can be used to compute the peak gain G_p.

$$|H(\omega_p)|^2 = 0 \tag{3.51}$$

$$\Rightarrow \omega_p^2 = \omega_0^2 \cdot \frac{2Q^2 - 1}{2Q^2} \tag{3.52}$$

$$G_p^2 = |H(\omega)|^2\Big|_{\omega=\omega_p} = \frac{4Q^4}{4Q^2 - 1} \tag{3.53}$$

Using the analog peak location and gain, we can solve for B by constraining the digital peak gain to be equal to G_p.

$$|H_d(\omega = \omega_0)|^2 = G_p^2 \tag{3.54}$$

$$\Rightarrow B^2 = \frac{\omega_p^2 \omega_d^2 G_p^2 + Q_d^2 G_p^2 (\omega_p^2 - \omega_d^2)^2 - Q_d^2 (\omega_d^2 - G_1 \omega_p^2)^2}{\omega_p^2 \omega_d^2} \tag{3.55}$$

We can then find ω_d by setting the digital magnitude derivative equal to zero and substituting in the above equation for B.

$$\frac{\delta}{\delta \omega} |H(\omega)|^2 \Bigg|_{\omega=\omega_0, B} = 0 \tag{3.56}$$

$$\Rightarrow \omega_d^2 = \omega_p^2 \sqrt{\frac{G_p^2 - G_1^2}{G_p^2 - G_0^2}} \tag{3.57}$$

Lastly, to solve for Q_d, we can select an arbitrary gain between unity and the peak gain to constrain the width of the digital peak. As in the previous filters, we will choose $G_B = \sqrt{G_p}$. To find the frequency corresponding to G_B, we can set the analog magnitude response equal to G_B. Note that due to the symmetry of the filter peak, there are two points corresponding to G_B. We will choose the lower one, as that one is guaranteed to be below the cutoff frequency ω_0 and therefore below the Nyquist frequency.

$$|H(s)|^2 = G_B^2 = G_p \tag{3.58}$$

$$\Rightarrow \omega_B^2 = \omega_0^2 \cdot \frac{G_p(2Q^2-1) - \sqrt{G_p(G_p - 4G_pQ^2 + 4Q^4)}}{2G_pQ^2} \tag{3.59}$$

Similarly, we can set the digital magnitude response equal to G_B and also substitute $\omega \leftarrow \omega_B$ and B with Equation 3.55.

$$|H_d(\omega = \omega_B)|^2\Big|_{\omega=\omega_{0,B}} = G_p \tag{3.60}$$

$$\Rightarrow Q_d^2 = G_p\omega_B^2\omega_d^2\omega_p^2(1-G_p)/(G_1^2\omega_B^4\omega_p^2 - G_1^2\omega_B^2\omega_p^4 + G_p^2\omega_B^2\omega_d^4 - 2G_p^2\omega_B^2\omega_d^2\omega_p^2 + \\ + G_p^2\omega_B^2\omega_p^4 - G_p\omega_B^4\omega_p^2 + 2G_p\omega_B^2\omega_d^2\omega_p^2 - G_p\omega_d^4\omega_p^2 - \omega_B^2\omega_d^4 + \omega_d^4\omega_p^2) \tag{3.61}$$

3.5.2 Solving with no peak

When $Q \leq 1/sqrt2$, a peak is no longer present in the magnitude response, so the above solution will not work. We can use an alternative approach to constrain the parameters. To begin, we can simply set $\omega_d = \omega_0$, as there is no peak which we would otherwise need to correctly position. We can solve for B and Q_d by matching gain at various points. In particular, we can compute B by setting the digital gain at ω_d to be equal to the analog gain at ω_0, which we found earlier to be Q.

$$|H_d(\omega = \omega_d)|^2 = Q^2 \tag{3.62}$$

$$\Rightarrow B^2 = Q^2 - Q_d^2(G_1 - 1)^2 \tag{3.63}$$

To find Q_d, we need to select a second point at which to match the magnitude response. By finding the point at which the analog magnitude is equal to $\sqrt{Q}$, we can verify that this frequency will always be below the cutoff frequency ω_0 and thus be below the Nyquist frequency.

We first find this half-gain analog frequency ω_B and then use it to constrain the digital gain.

$$\left.|H(\omega)|^2\right|_{\omega=\omega_p} = Q \tag{3.64}$$

$$\Rightarrow \omega_B^2 = \omega_0^2 \cdot \frac{2Q^2 - 1 + \sqrt{4Q^3 - 4Q^2 + 1}}{2Q^2} \tag{3.65}$$

$$\left.|H_d(\omega)|^2\right|_{\omega=\omega_{B,B}} = Q \tag{3.66}$$

$$Q_d^2 = Q \cdot \frac{\omega_B^2 \omega_d^2 (1-Q)}{(\omega_B^2 - \omega_d^2)(G_1^2 \omega_B^2 - Q\omega_B^2 + Q\omega_d^2 - \omega_d^2)} \tag{3.67}$$

3.5.3 Combining the solutions

To mitigate any discontinuities present between the two filter parameter derivations, we can interpolate between them. It is important to interpolate only when $Q > 1/\sqrt{2}$, as otherwise the first solution, which relies on a peak being present, will produce undefined results. The resulting digital parametric low-pass filter response can be seen in Figure 3.14. We have successfully designed a digital filter whose magnitude response closely matches its analog counterpart without increasing the filter order or sacrificing latency and stability (Figure 3.15). Figures 3.16 and 3.17 demonstrate that the filter maintains its desired response for small and large values of Q, respectively.

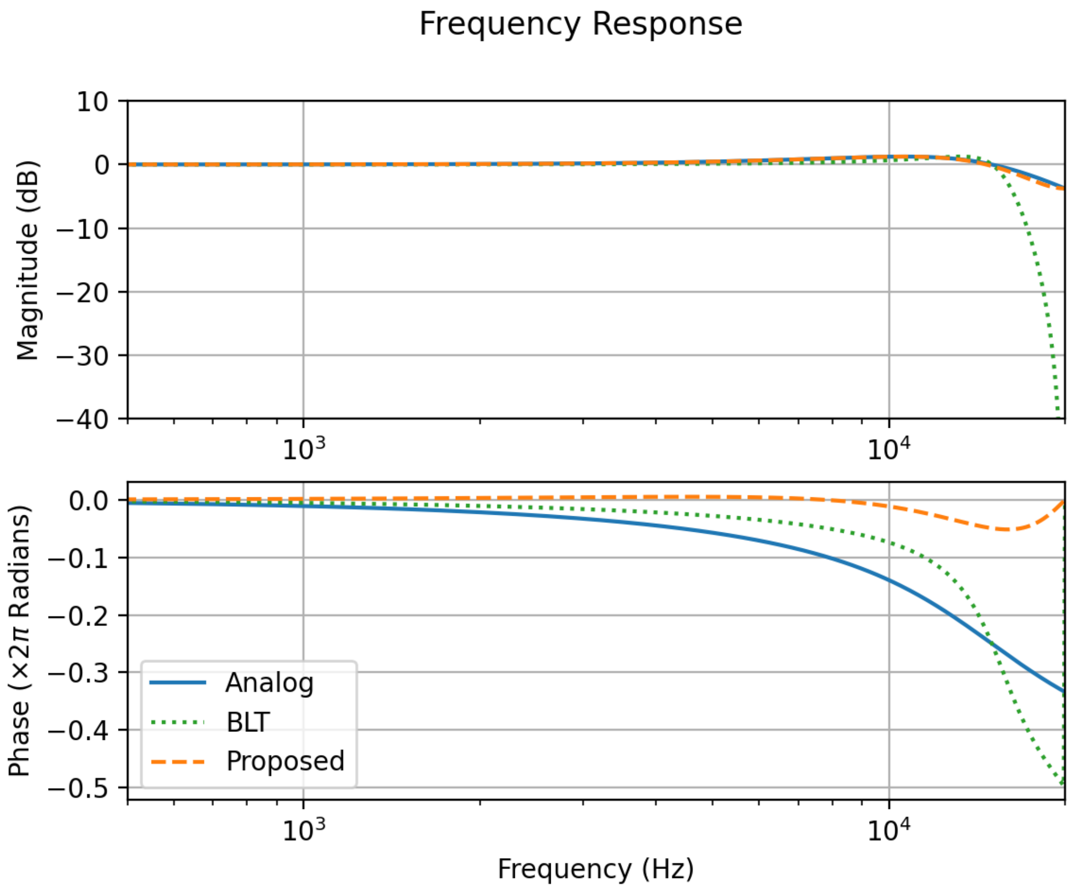

Figure 3.14. Proposed method to derive the low-pass filter ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 1$).

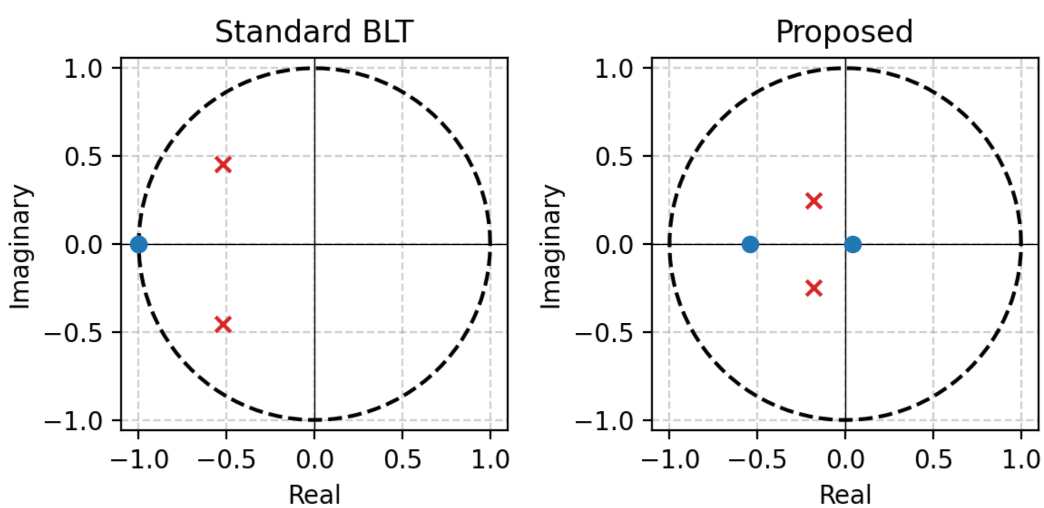

Figure 3.15. Pole-zero diagrams of low-pass filters ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 1$).

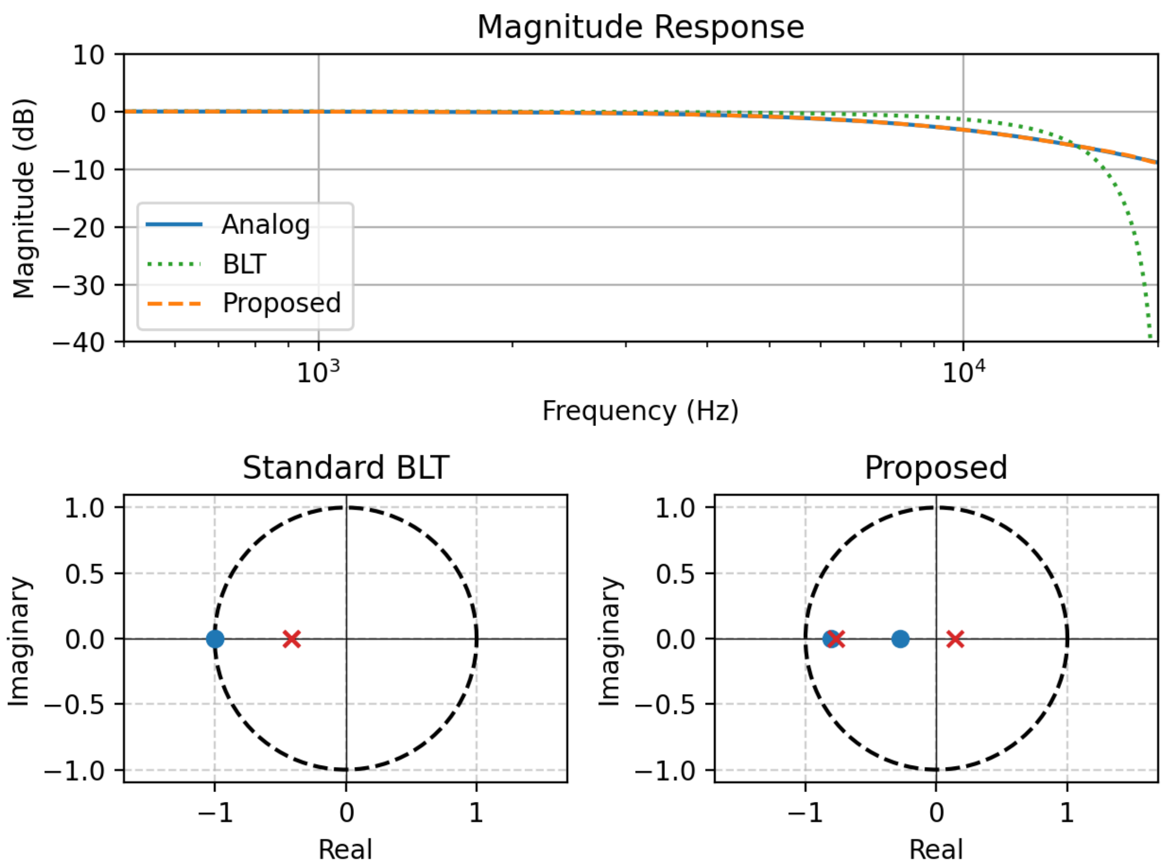

Figure 3.16. Low-pass with small Q ($F_s = 40$kHz, $f_0 = 15$kHz, $Q = 0.5$).

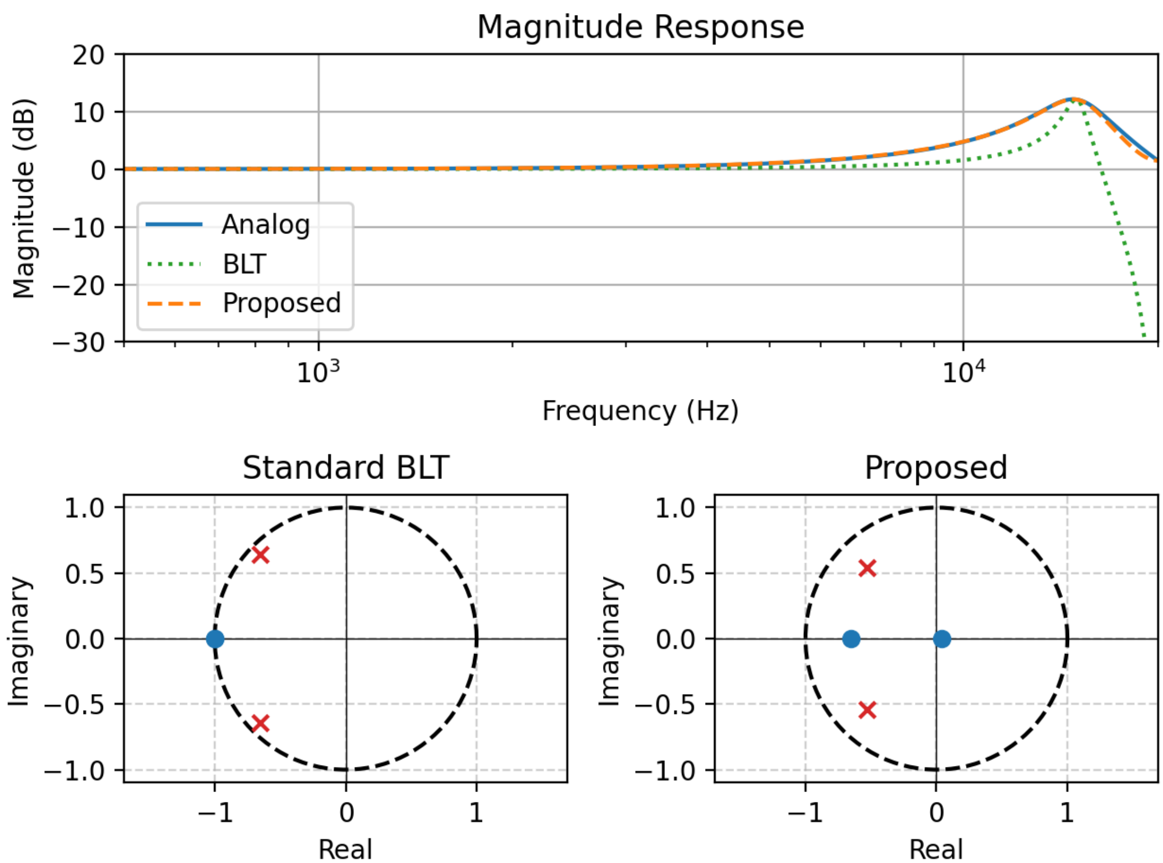

Figure 3.17. Low-pass with large Q ($F_s = 40\text{kHz}$, $f_0 = 15\text{kHz}$, $Q = 4$).

3.6 Shelving Filters

Of all the parametric filters discussed thus far, the shelving filters proved to be the most difficult. For instance, the high-shelf filter's transfer function [2] is shown in Equation 3.68. Whereas previous filters had magnitude responses with at most a single peak, the shelving filters can have up to two. In addition, since the cutoff frequency corresponds to the center of the transition band, the upper peak could at times be above the Nyquist frequency.

$$H(s) = A\frac{As^2 + \omega_0\sqrt{A}/Q + \omega_0^2}{s^2 + \omega_0\sqrt{A}/Q + \omega_0^2 A}, \quad \text{where } A = \sqrt{g} \tag{3.68}$$

There are three important characteristics of shelving filters: the transition band, the shelving band, and sometimes the peaks. As a result, it becomes very difficult to define a second-order digital filter that closely matches its analog counterpart at these key points. Unfortunately, due to time limitations, we have not been able to construct digital shelving filters which mitigate the frequency warping issue. However, that does not imply designing such a filter is impossible; more research is needed to determine a system of constraints which will allow the derivation of such shelving filters.

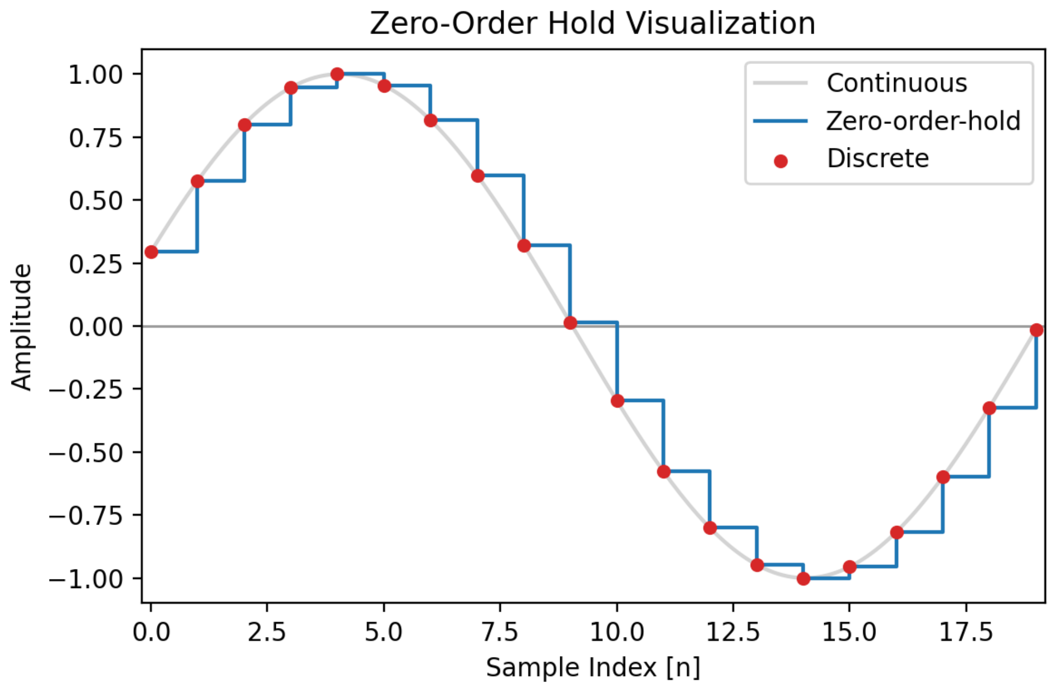

Figure 3.18. Zero-order-hold of a discrete time series in an ideal DAC.

3.7 Pre-emphasis Filter

Although not part of the parametric family of filters, pre-emphasis filters are an important piece of the digital audio puzzle. Before we are able to hear digital audio signals, they must first pass through a digital-to-analog converter (DAC). Many modern DACs use a technique referred to as "sample-and-hold" or "zero-order-hold". When this type of DAC receives a digital sample, it will output a voltage proportional to the sample amplitude. The DAC will then maintain the voltage level until the next sample comes in, demonstrated in Figure 3.18. This, of course, impacts the frequency response of the signal.

The result is identical to convolving the input sample impulses with a rectangle function, where the rectangle width is $1/F_s$ seconds in length. With this perspective, we can compute the effect this convolution has on the frequency response of the signal. Since the Fourier Transform of the rectangle function is $\text{sinc}(ax) = \sin(ax)/(ax)$ where $a = \pi/F_s$, we can deduce that the magnitude response will be low-passed by the sinc function, as shown in Figure 3.19.

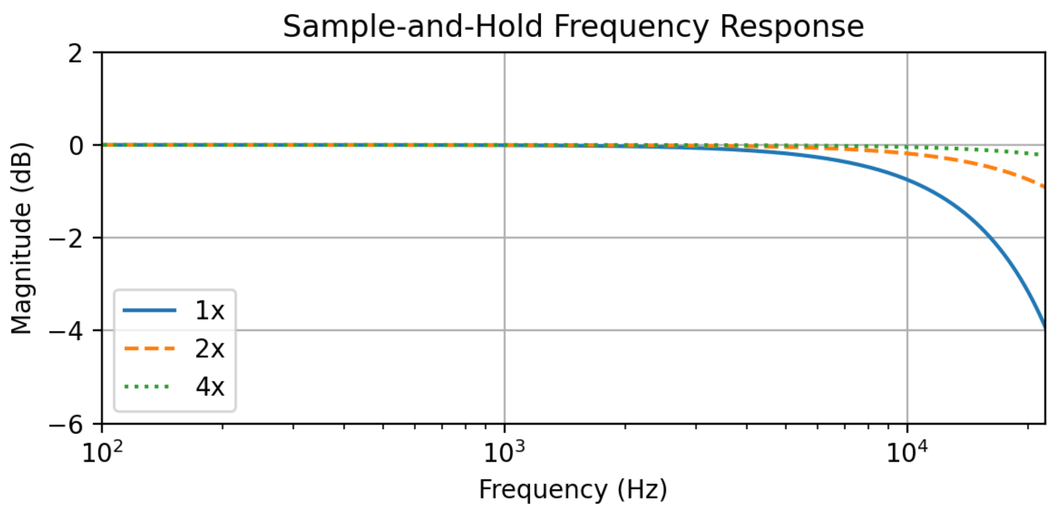

Figure 3.19. Frequency response of zero-order-hold at different upsampling levels ($F_s =$ 44.1kHz).

To combat the unwanted attenuation of the high frequency, DACs often employ a pre-emphasis filter. Ideally, the pre-emphasis filter, call it $H_p(s)$, will have a magnitude response which is the inverse of the sinc function, as shown in Figure 3.20, resulting in Equation 3.69.

$$|H_p(f)| = \frac{1}{\text{sinc}(af)} = \frac{af}{\sin(af)} \tag{3.69}$$

In DACs, such as the Texas Instruments DAC5688, a digital 9-tap FIR filter is used [8]. However, using the techniques presented, we can derive a digital IIR pre-emphasis filter with fewer multiplies per sample. Important to keep in mind is that DACs often up-sample the incoming audio signals by 2 or 4 times. This allows for cheaper analog filters to be used when anti-aliasing the DAC output. As a result, the attenuation caused by the sample-and-hold process is also reduced. Thus, even lower order pre-emphasis filters can be used. For instance, if up-sampling, even first-order digital IIR pre-emphasis filters match the ideal magnitude response closely.

Furthermore, unlike the parametric filters, the pre-emphasis filter coefficients can be precomputed, as the filter only needs adjusting if the up-sampling amount changes. Thus, slower

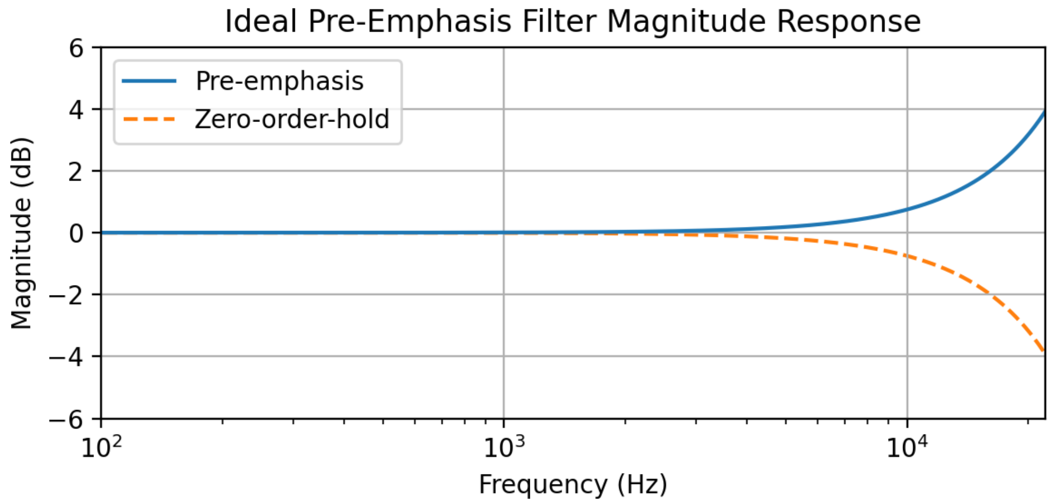

Figure 3.20. Magnitude response of the ideal pre-emphasis filter ($F_s = 44.1$kHz, no upsampling).

iterative techniques, such as gradient descent, can be utilized. The pre-emphasis filters presented below are derived using the general digital biquadratic IIR filter in Equation 3.1. We have determined these parameter values via manual tuning.

Table 3.1 displays the intermediate digital parameter values to achieve pre-emphasis filters, depending on the amount of digital interpolation in the system. The equations found at the start of Chapter 3 can be used to compute the filter coefficients.

Table 3.1. Intermediate digital parameters for pre-emphasis filters accounting for various levels of interpolation.

Upsampling	G_1	B	ω_d	Q_d
1x	1.508	1.259	2.4689	0.44
2x	1.1032	1.0594	2.4142	0.424
4x	1.0244	1.0145	2.4142	0.42

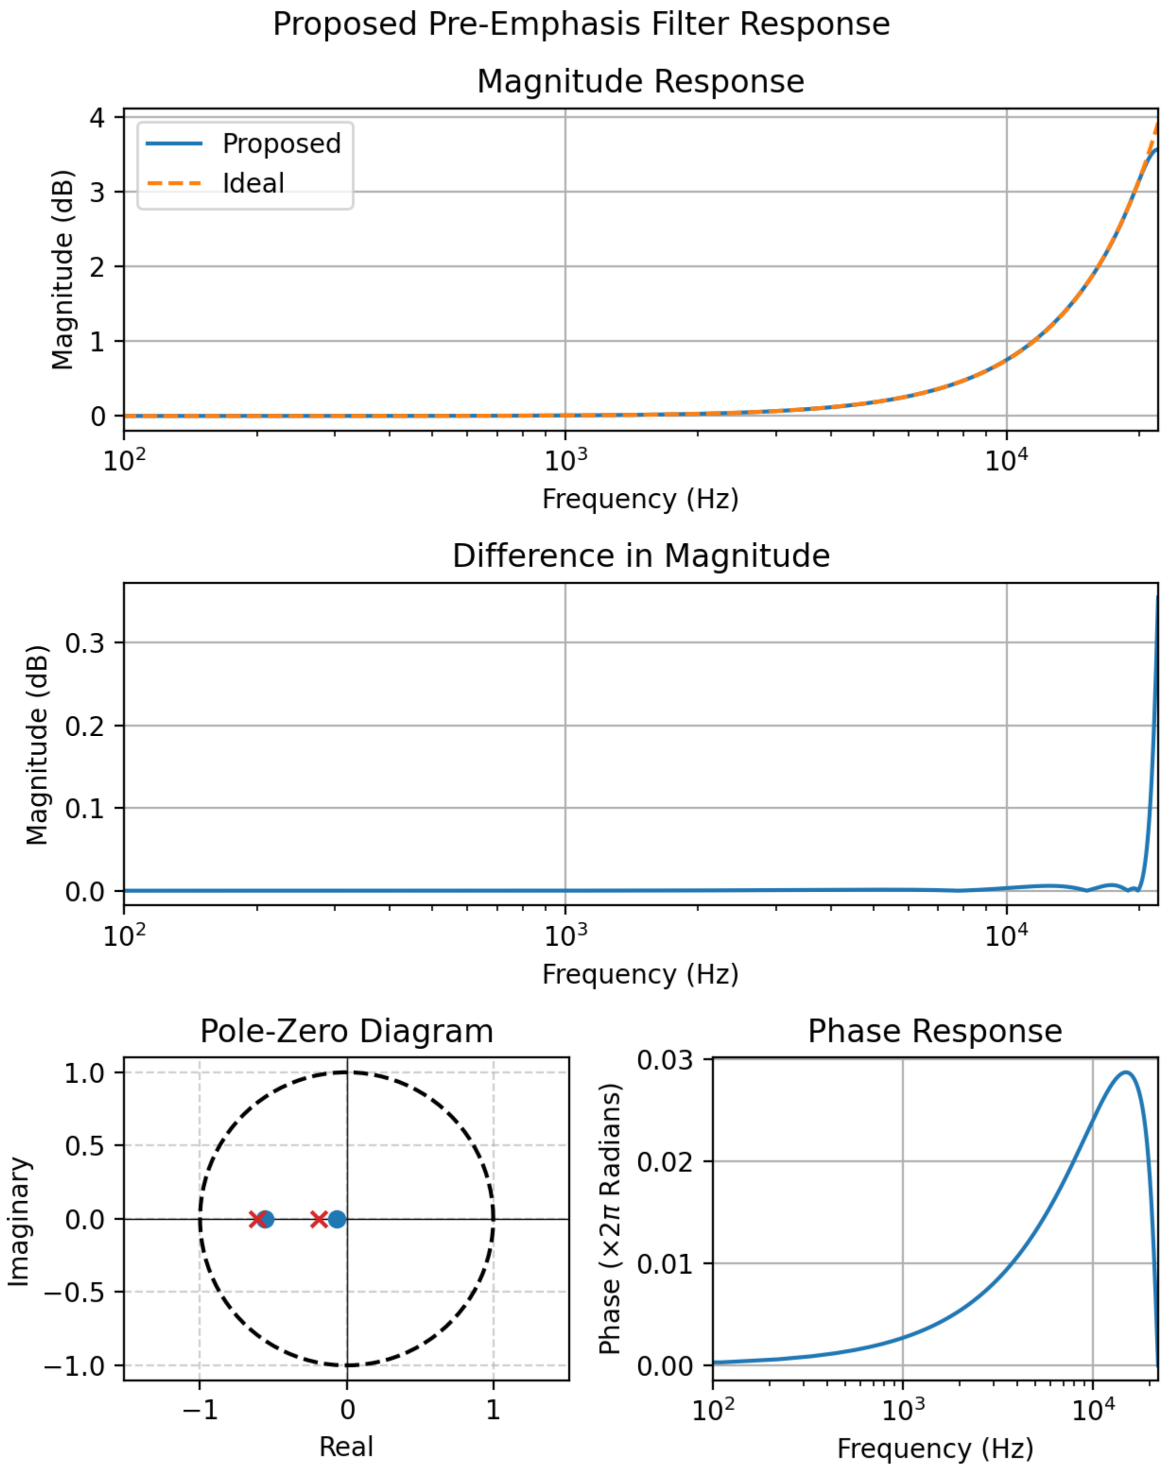

Figure 3.21. Frequency response of the proposed pre-emphasis filter ($F_s = 44.1$kHz, no upsampling).

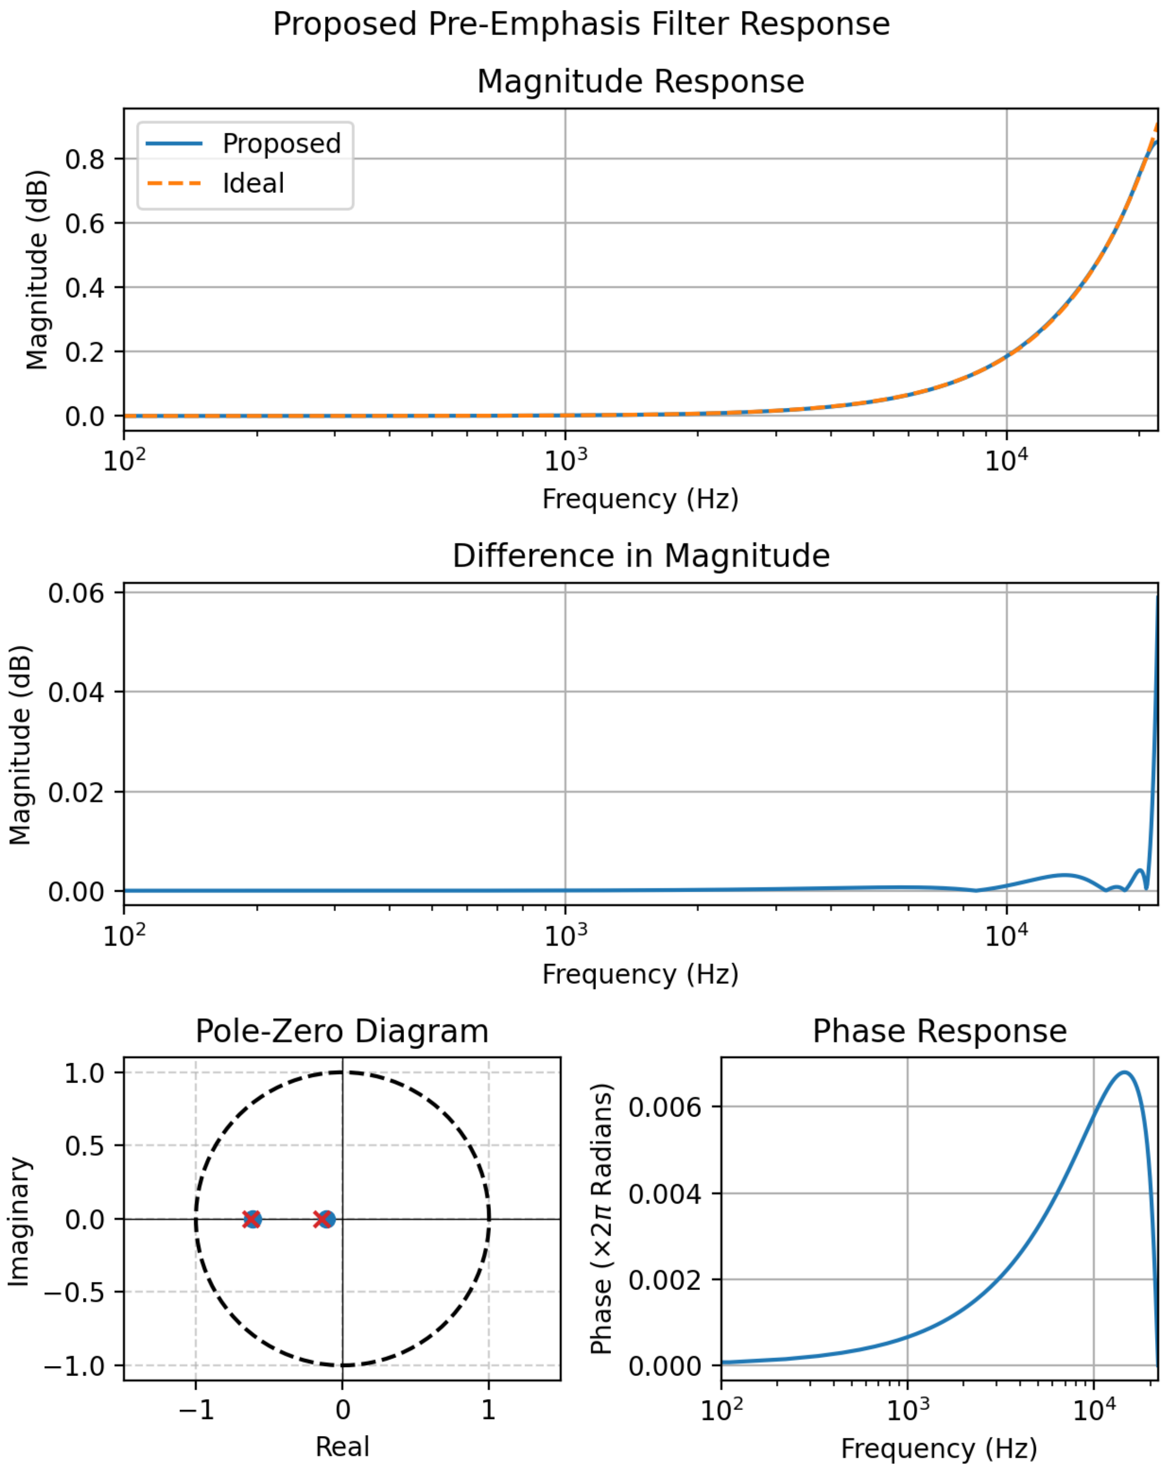

Figure 3.22. Frequency response of the proposed pre-emphasis filter ($F_s = 44.1$kHz, 2x upsampling).

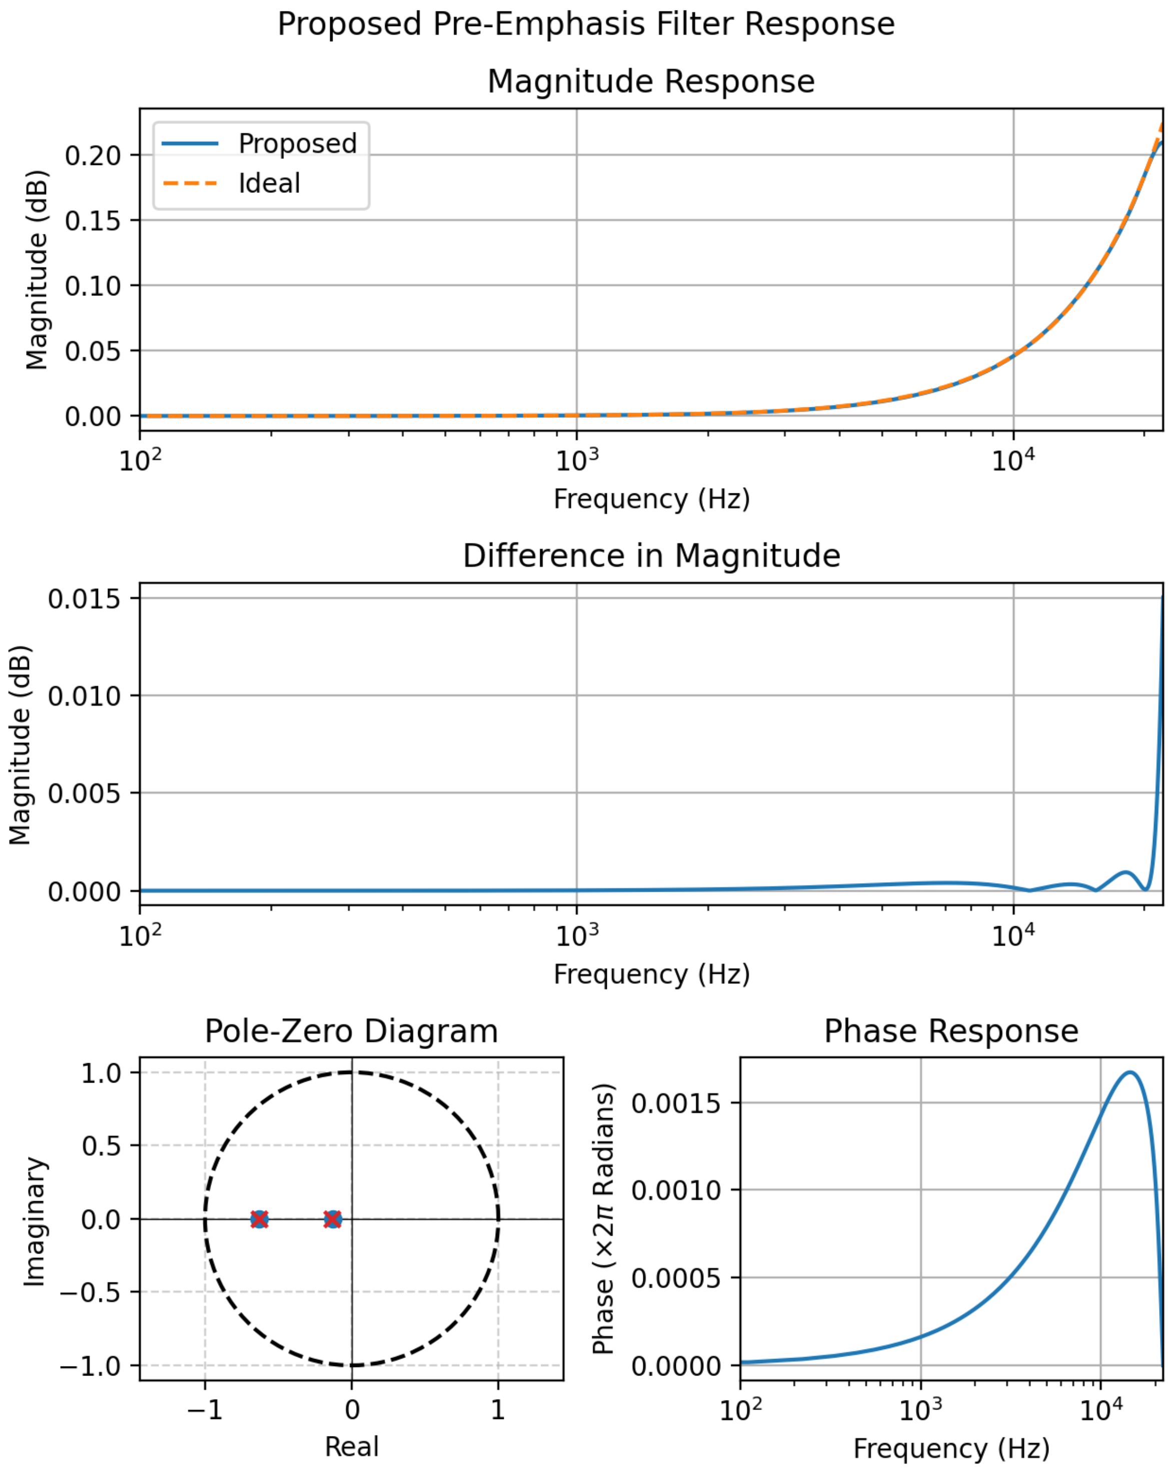

Figure 3.23. Frequency response of the proposed pre-emphasis filter ($F_s = 44.1$kHz, 4x upsampling).

Chapter 4

Conclusion

4.1 Comparison

To further demonstrate the accuracy of the proposed technique, Figure 4.1 compares the magnitude response of the various peaking filter designs discussed. By computing the difference between the original analog and various digital magnitude responses, we can evaluate the accuracy of each digital approach. This metric confirms that our proposed approach provides a magnitude response which most closely matches the analog response. We have also provided pole-zero diagrams of each digital filter in Figure 4.2.

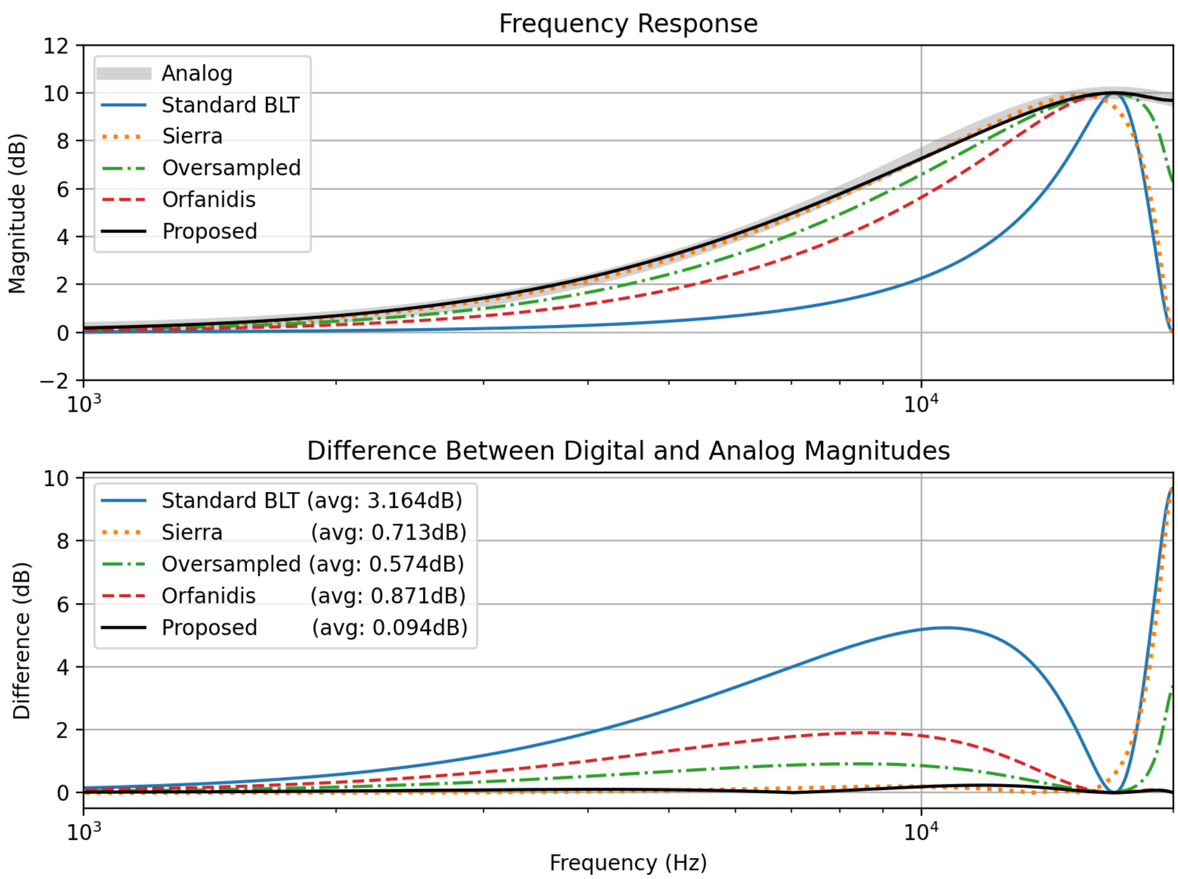

Figure 4.1. Comparison of all parametric filters discussed ($F_s = 40$kHz, $f_0 = 17$kHz, $Q = 0.5$, $g = 10$dB).

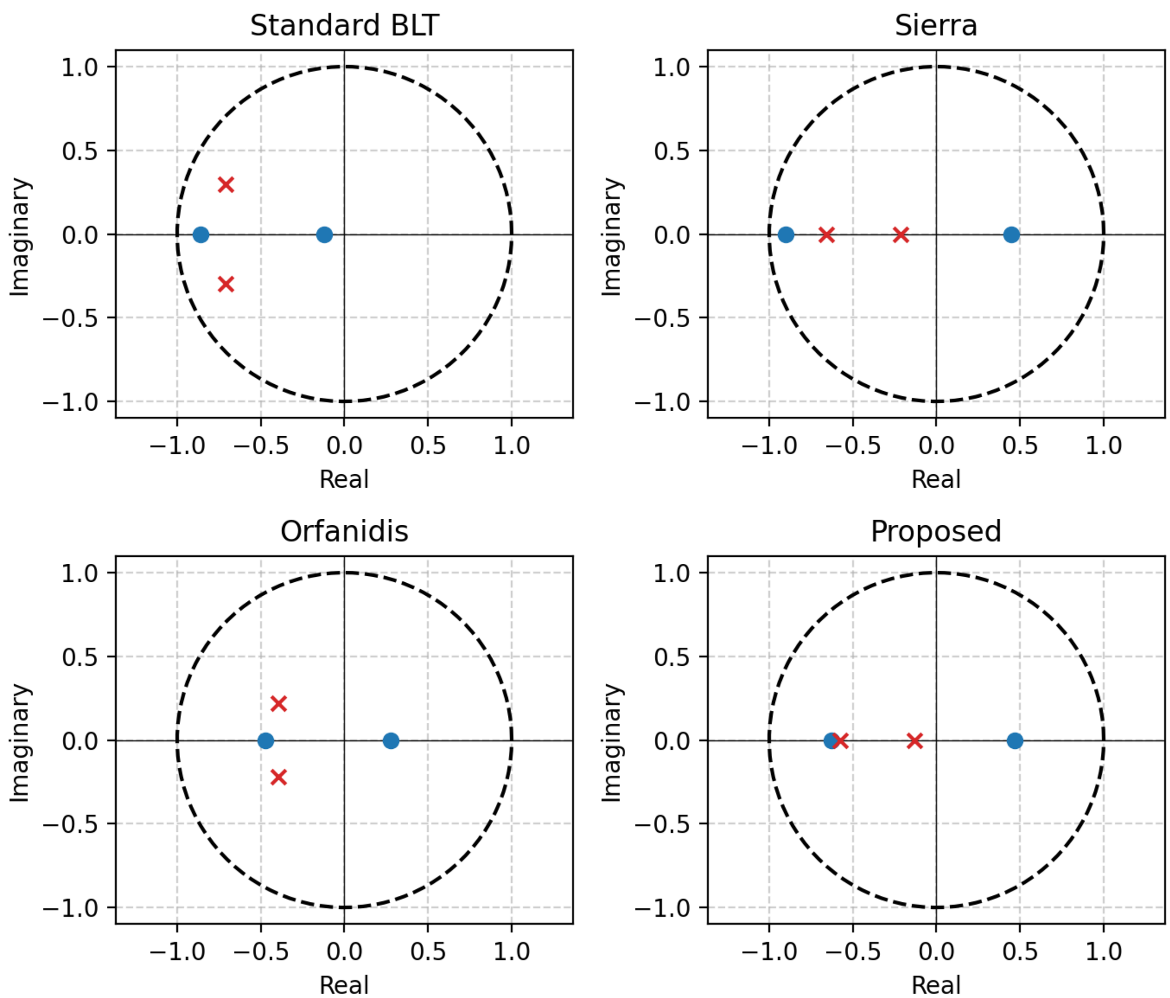

Figure 4.2. Pole-zero diagrams of all parametric filters discussed ($F_s = 40$kHz, $f_0 = 17$kHz, $Q = 0.5$, $g = 10$dB).

We can also compare the proposed approach to existing commercial software packages which also claim to mitigate the frequency warping effect. For this, we have chosen FabFilter's Pro-Q 4; not only is it widely used throughout the audio industry, but its marketing also asserts "Pro-Q matches the magnitude response of analog EQing as closely as possible". Figures 4.3, 4.4, 4.5, 4.6, and 4.7 compare each of the five parametric filters derived in Chapter 3. The impulse responses are measured with Pro-Q 4 set to zero-latency mode and all other settings left at their default configurations.

It is important to note that Pro-Q 4's definition of the Q parameter differs from the one displayed in the Audio EQ Cookbook. This is most noticeable when attempting to achieve Butterworth filters: $Q = 1/\sqrt{2}$ for the standard definition and $Q = 1$ for Pro-Q 4. We can conclude that the engineers at FabFilter have skewed the values of Q. In order to create accurate comparisons, we must account for these differences. We have managed this by manually tuning Pro-Q 4's Q value until the magnitude responses closely match the proposed filter responses. Importantly, this matching is done at the low cutoff frequency of $f_0 = 1\text{kHz}$ in order to ensure frequency warping is not present. For instance, when the proposed filter quality is set to $Q = 1.2$, that roughly corresponds to $Q = 1.696$ in Pro-Q 4.

We can confidently say that Pro-Q 4's zero-latency mode is using a similar technique to the one proposed; we observe that our proposed and Pro-Q 4 phase responses are very similar and both deviate from their analog counterparts. The similarity in phase response allows us to speculate that oversampling is not used, nor is a higher order filter utilized, as either technique could closely match the digital phase response to the analog counterpart [10].

While the proposed and Pro-Q 4 magnitude responses are also very similar, notice how in most filters, the gain at Nyquist is not matched. We speculate that FabFilter's reason for doing this is to sacrifice the magnitude response accuracy at the highest frequencies in exchange for a more accurate response at lower, more audible frequencies.

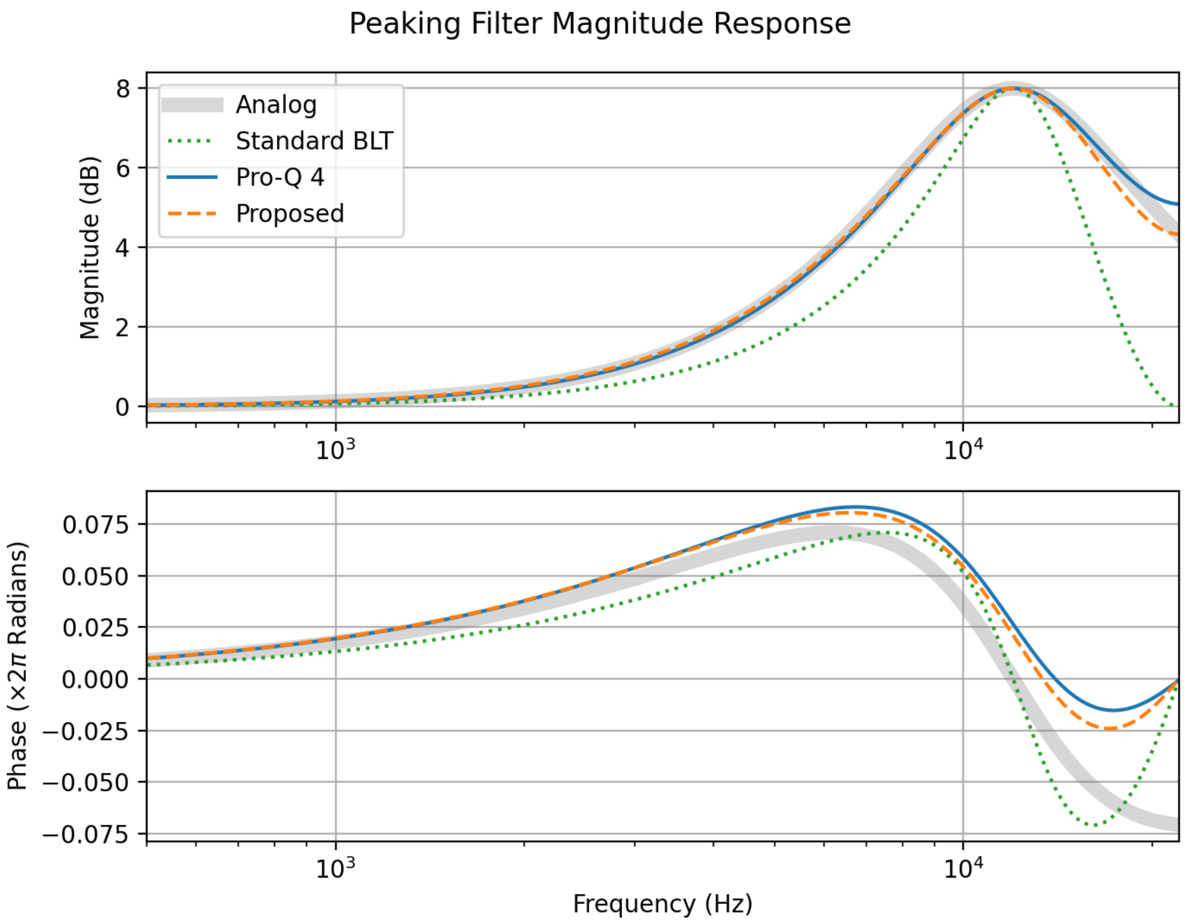

Figure 4.3. Comparison with FabFilter Pro-Q 4's peaking filter ($F_s = 44.1$kHz, $f_0 = 12$kHz, $Q = 1/\sqrt{2}$, $g = 8$dB).

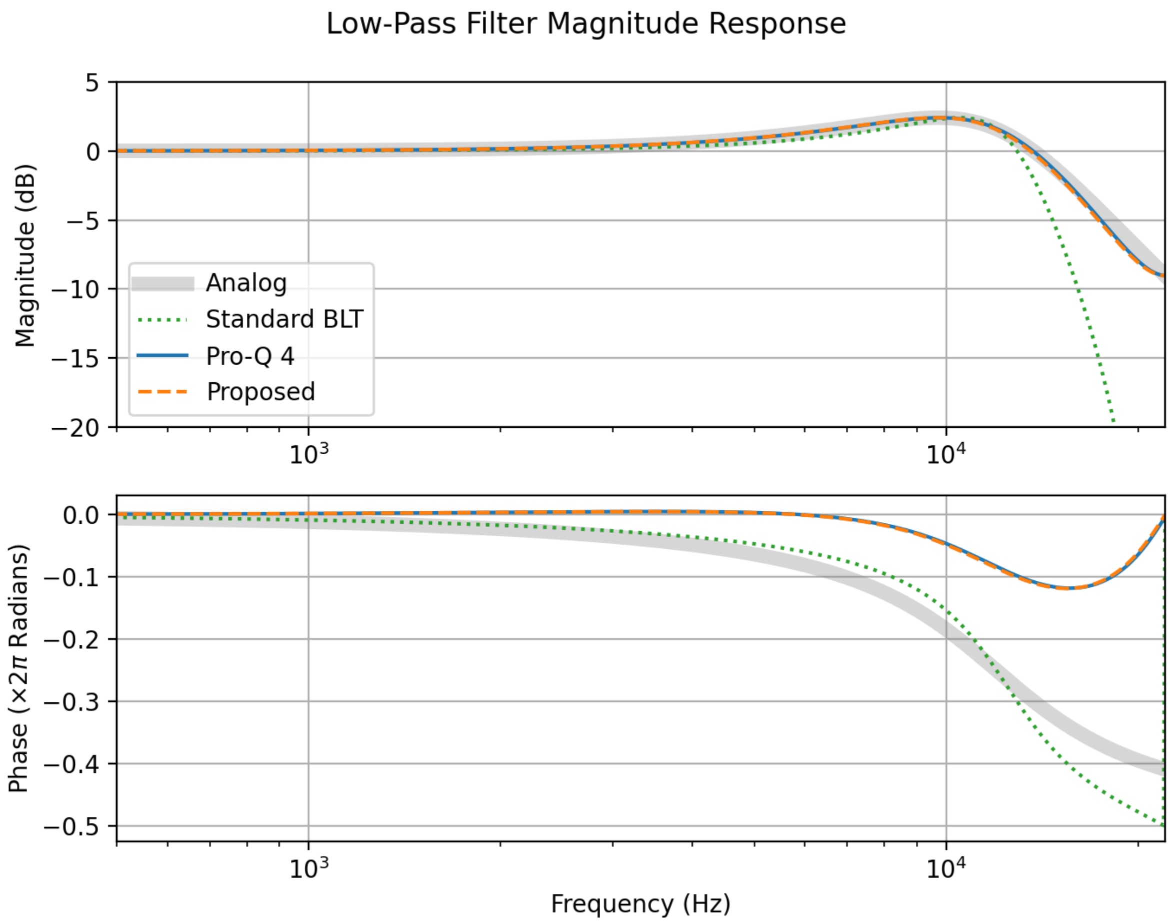

Figure 4.4. Comparison with FabFilter Pro-Q 4's low-pass filter ($F_s = 44.1$kHz, $f_0 = 12$kHz, $Q = 1.2$).

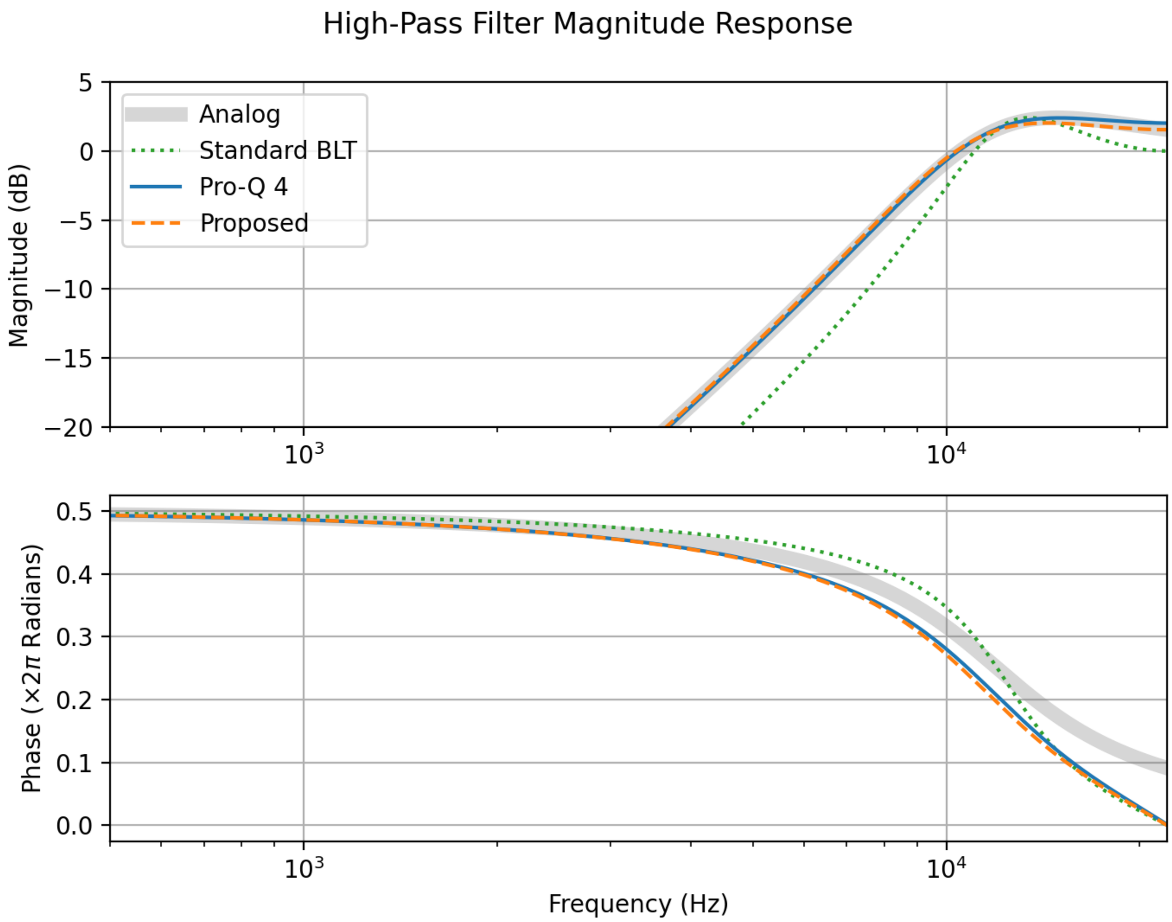

Figure 4.5. Comparison with FabFilter Pro-Q 4's high-pass filter ($F_s = 44.1$kHz, $f_0 = 12$kHz, $Q = 1.2$).

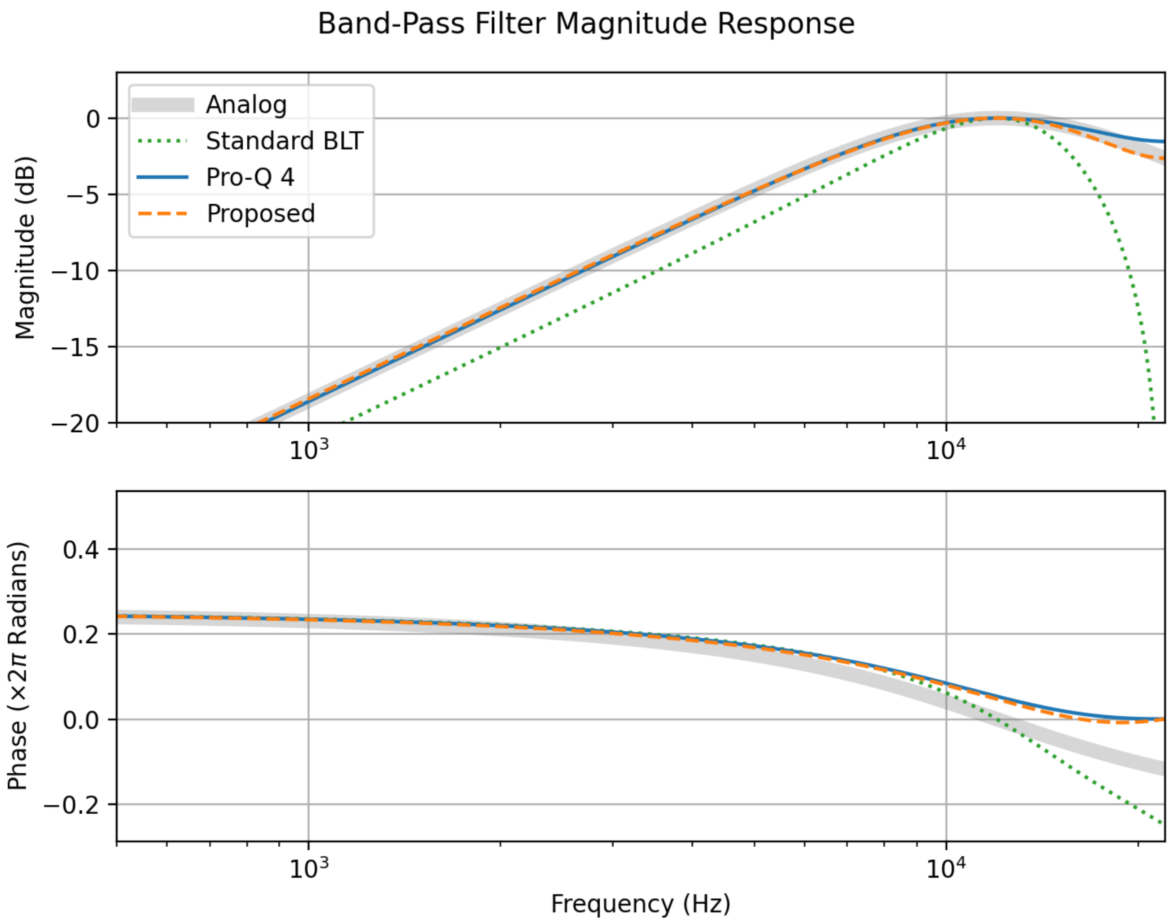

Figure 4.6. Comparison with FabFilter Pro-Q 4's band-pass filter ($F_s = 44.1$kHz, $f_0 = 12$kHz, $Q = 1/\sqrt{2}$).

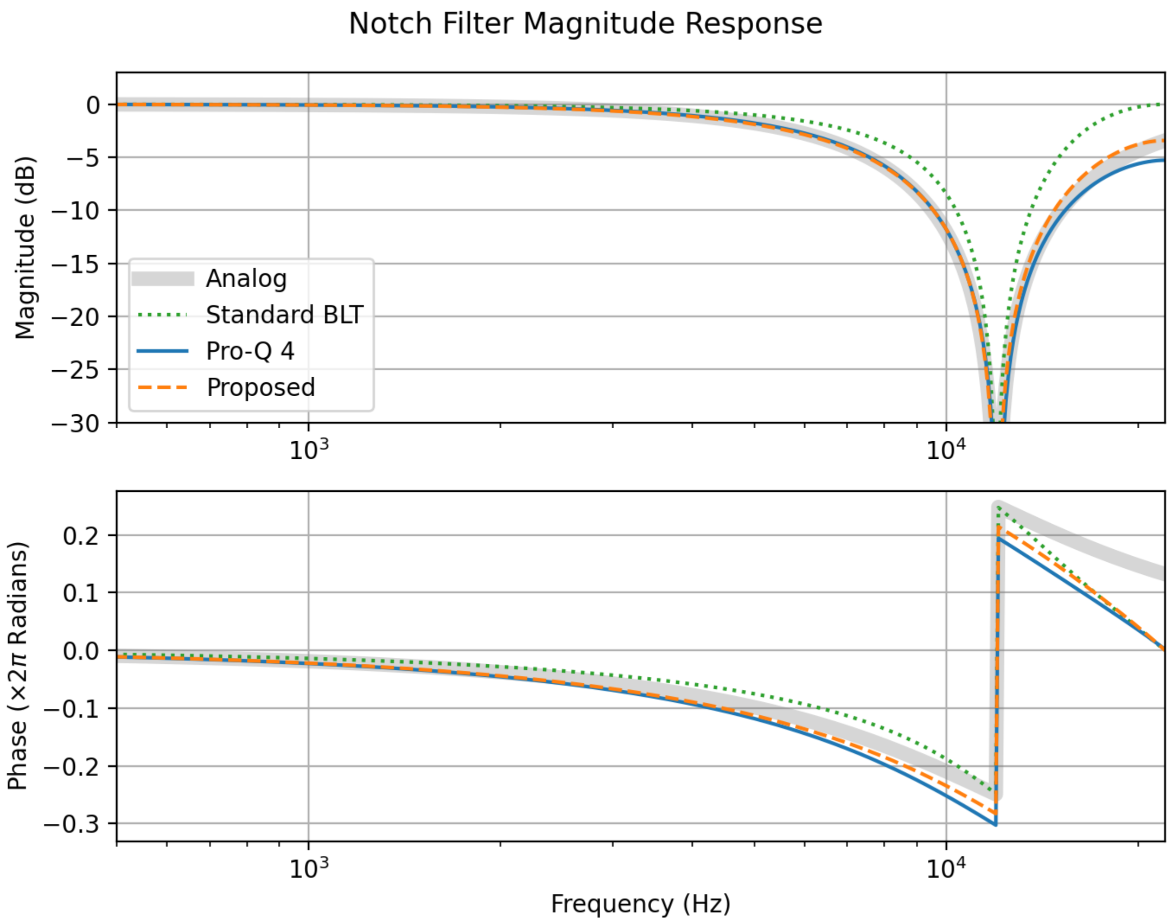

Figure 4.7. Comparison with FabFilter Pro-Q 4's notch filter ($F_s = 44.1$kHz, $f_0 = 12$kHz, $Q = 1/\sqrt{2}$).

4.2 Further Work

As discussed in Section 3.6, we did not have the time to derive shelving filters with the proposed approach. More research is necessary to determine the right set of constraints which would allow us to solve for such filters.

Similarly to the standard BLT-derived parametric filters, the proposed filters become unstable when cutoff frequencies ω_0 above the Nyquist frequency are selected. In this area, Sierra's [11] approach, as discussed in Section 2.3, is superior. Further research is required to allow for this kind of cutoff frequency flexibility in conjunction with the proposed approach.

Bibliography

[1] Andreas Antoniou. Bilinear-transformation method. In *Digital Filters: Analysis, Design, and Signal Processing Applications*, chapter 12.6. McGraw-Hill Education, New York, 1st edition, 2018.

[2] Robert Bristow-Johnson. Audio eq cookbook. Online, 2024. Available at https://webaudio.github.io/Audio-EQ-Cookbook/audio-eq-cookbook.html.

[3] A. G. Constantinides. Digital signal processing. In *UK 7th Conference: Digital Signal Processing (DSP)*, London, UK, September 1992. Audio Engineering Society. Paper Number: DSP-08.

[4] Champ Darabundit and Russell Wedelich. Generalized digital second order systems beyond nyquist frequency. In *Audio Engineering Society Convention 149*, pages 1–5, New York, NY, October 2020. Audio Engineering Society. Engineering Brief 631.

[5] F. Harris and E. Brooking. A versatile parametric filter using an imbedded all-pass sub-filter to independently adjust bandwidth, center frequency, and boost or cut . In *Conference Record of the Twenty-Sixth Asilomar Conference on Signals, Systems & Computers*, pages 269,270,271,272,273, Los Alamitos, CA, USA, October 1992. IEEE Computer Society.

[6] Alan V. Oppenheim, Alan S. Willsky, and S. Hamid Nawab. *Signals and Systems*. Prentice Hall, Upper Saddle River, New Jersey, 2nd edition, 1996.

[7] Sophoncles J. Orfanidis. Digital parametric equalizer design with prescribed nyquist-frequency gain. *Journal of The Audio Engineering Society*, 45:444–455, 1997.

[8] Chris Pearson. High speed, digital to analog converters basics. Application Report SLAA523A, Texas Instruments, October 2012. Rev. A.

[9] Dale Purves, George J. Augustine, David Fitzpatrick, Lawrence C. Katz, Anthony-Samuel LaMantia, James O. McNamara, and S. Mark Williams, editors. *Neuroscience*. Sinauer Associates, Sunderland, MA, 2nd edition, 2001. The Audible Spectrum.

[10] Simo Sarkka and Antti Huovilainen. Accurate discretization of analog audio filters with application to parametric equalizer design. *IEEE Transactions on Audio, Speech, and Language Processing*, 19(8):2486–2493, 2011.

[11] Juan Sierra. Digital parametric filters beyond nyquist frequency. *Journal of the Audio Engineering Society*, (10224), October 2019.

[12] J.O. Smith. *Introduction to Digital Filters: With Audio Applications*. Music signal processing series. W3K, 2007.

[13] Anastasiia Sokolova, Dwaipayan Sengupta, Matthew Hunt, Ramesh Gupta, Baris Aksanli, Fredric Harris, and Chandrasekhar Garudadri. Real-time multirate multiband amplification for hearing aids. *IEEE Access*, 10:54301–54312, 2022. Epub 2022 May 20.

[14] Reinaldo Valenzuela and A. Constantinides. Digital signal processing schemes for efficient interpolation and decimation. *Electronic Circuits and Systems, IEE Proceedings G*, 130:225 – 235, 01 1984.

[15] Mateusz Żuk and Joshua D. Reiss. Aural phase distortion detection. In *Proceedings of the 26th International Conference on Digital Audio Effects (DAFx)*, pages 1–8, Copenhagen, Denmark, September 2023.