

Lecture 8

DT Filter Design: IIR Filters

Reading: Section 7.1 in Oppenheim, Schafer & Buck (OSB).

In the last lecture we studied various forms of filter realizations. Today we will take one step back, focusing our attention on determining the actual transfer functions to be realized. Given specifications of the desired properties of the system, we approximate the specifications using a causal discrete time system. In this lecture, we will show how IIR systems can be approximated by rational functions of z . Next time we will look at FIR system approximation using polynomials of z .

So far, we have tended to make reference only to ideal filters. In practical designs, ideal systems can not be realized exactly, since the impulse responses are infinitely long and non-causal. The idea is to approximate the desired frequency response within certain error tolerances. As such, specifications are usually given as tolerance levels in different frequency bands in the range $0 \leq \omega \leq \pi$. OSB Figure 7.2 gives a sample tolerance scheme. The transition band is a “don’t care” region in which the filter gain can be any finite value. Notice that no requirements have been specified for the phase response of the system. Typical filter design procedures focus only on magnitude approximation:

$$H_{\text{ideal}}(e^{j\omega}) = |H(e^{j\omega})|e^{j\theta(\omega)} \quad \rightarrow \quad |H_{\text{eff}}(e^{j\omega})| \approx |H_{\text{ideal}}(e^{j\omega})| = |H(e^{j\omega})|.$$

Nonetheless, specifications can involve both magnitude and phase (or group delay). Such generalized approximation is a harder problem, but may be desired in specific applications. In particular, integer or fractional delays can only be achieved with FIR filters.

Most of our following discussions will be phrased for piecewise constant LPF’s. Keep in mind, however, that much applies more generally, since specific transformations can convert LPF’s to HP, BP, or notch filters.

IIR vs. FIR

Given a set of specifications, first we need to decide if the desired filter should be IIR or FIR. The following table summarizes different factors that could be considered when making this decision:

	IIR Filters	FIR Filters
Phase (grp delay)	difficult to control, no particular techniques available	linear phase always possible
Stability	can be unstable, can have limit cycles	always stable, no limit cycles
Order	less	more
History	derived from analog filters	no analog history
Others		polyphase implementation possible can always be made causal

- The only way to achieve integer or fractional constant delays is by using FIR filters.
- Limit cycles are instability of a particular type due to quantization, which is severely non-linear.
- The number of arithmetic operations needed per unit time is directly related to filter order. #MAD stands for number of multiplications and additions. When using this as a criterion for comparing different filters, one should pay attention to whether the #MAD is measured per input sample, per output sample, or per unit time (clock cycle). It is possible that an IIR of lower order actually requires more #MAD than an FIR of higher order, because FIR filters may be implemented using polyphase structures.
- Different conventions exist for specifying magnitude responses for IIR and FIR filters. In particular, IIR filter specs are normalized to $[1 - \Delta_1, 1]$ (dB) in the passband, and $[-\infty, \Delta_2]$ (dB) in the stopband; FIR filter specs are normalized to between $1 \pm \delta_1$ within the passband, and $\pm \delta_2$ in the stopband, where δ_1, δ_2 are given as decimals.

IIR Filter Design

Historically, digital IIR filters have been derived from their analog counterparts. There are several common types of analog filters: Butterworth which have maximally flat passbands in filters of the same order, Chebyshev type I which are equiripple in the passband, Chebyshev type II which are equiripple in the stopband, and Elliptic filters which are equiripple in both the passband and the stopband. The digital version of these can be obtained from analog designs through the bilinear transformation, discussed in detail in OSB Sections 7.1.2 and 7.1.3.

In short, the bilinear transformation is an algebraic mapping from the continuous frequency variable s to the discrete frequency variable z such that the imaginary axis in the s -plane corresponds to one revolution of the unit circle in the z -plane:

$$s \rightarrow \frac{1 - z^{-1}}{1 + z^{-1}} \quad \Rightarrow \quad j\Omega = j\frac{2}{T} \tan \frac{\omega}{2}, \quad \frac{\Omega}{\Omega_c} \rightarrow \frac{\tan \omega/2}{\tan \omega_c/2}$$

π in the digital frequency domain corresponds to infinity in the analog frequency domain. Note that the bilinear transformation is really only appropriate in mapping filters which approximate piecewise constant filters.

Butterworth Filters

- Two design parameters: order of the filter N , cutoff frequency ω_c . The squared magnitude response of a Butterworth filter has the following form:

$$|H_{BW}(e^{j\omega})|^2 = \frac{1}{1 + \left[\frac{\tan(\omega/2)}{\tan(\omega_c/2)} \right]^{2N}}$$

The tan function arises here because of the bilinear transformation. In the analog case, the transfer function is in the form of $1/1 + (\frac{\Omega}{\Omega_c})^{2N}$. See Appendix B.1 of OSB for analog Butterworth filter design techniques.

- The magnitude response of a Butterworth filter decreases monotonically with frequency. The order of the system can be estimated by examining the desired filter gain at the cut-off frequency.
- Integer round up may be required in estimating the order of the system. Specs are therefore often exceeded at the passband and stopband edges. In addition, the specs are often much exceeded in the stopband, with attenuation approaching zero ($-\infty$ dB) as frequency approaches π . This is because all zeros of the system are located at $z = -1$. Since the gain diminishes quickly as frequency increases. It is possible that a lower order filter exists such that it satisfies the given specifications, but does not exceed them as greatly as the Butterworth design.
- The following sets of figures are examples of Butterworth filter design. For an additional example, see OSB Example 7.4.

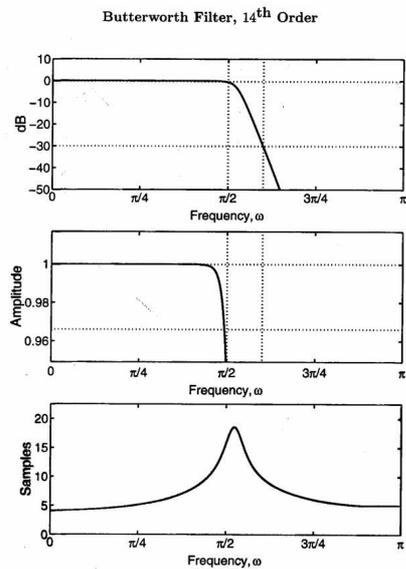


Figure 5: Magnitude Plot, Passband Detail, Group Delay

- Design Specifications:
1. Passband edge = 0.5π . ← Filter doesn't meet this spec
 2. Stopband edge = 0.6π .
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB. ← Filter doesn't meet this spec
 5. Maximum stopband gain = -30 dB.

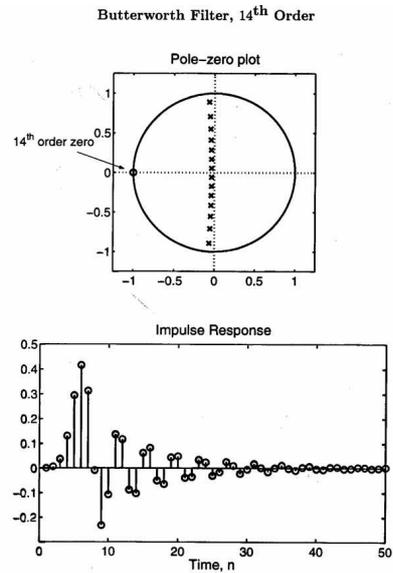


Figure 6: Pole-Zero Plot, Impulse Response

- Design Specifications:
1. Passband edge = 0.5π . ← Filter doesn't meet this spec
 2. Stopband edge = 0.6π .
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB. ← Filter doesn't meet this spec
 5. Maximum stopband gain = -30 dB.

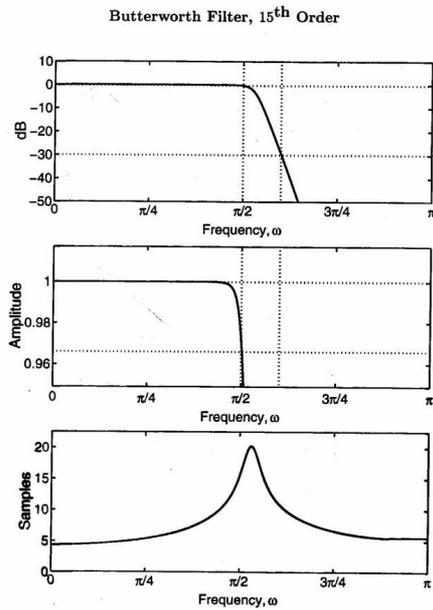


Figure 3: Magnitude Plot, Passband Detail, Group Delay

- Design Specifications:
1. Passband edge = 0.5π .
 2. Stopband edge = 0.6π .
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB.
 5. Maximum stopband gain = -30 dB.

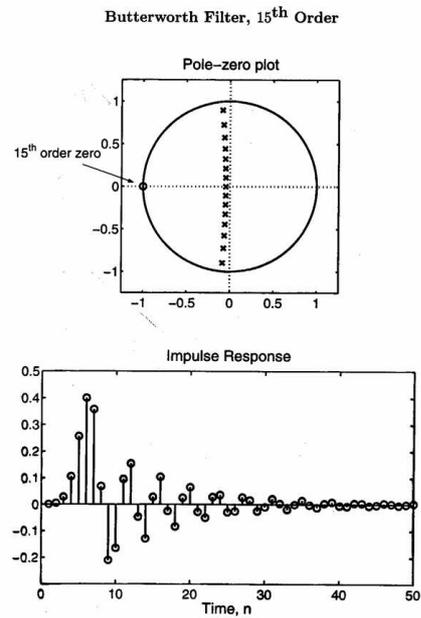
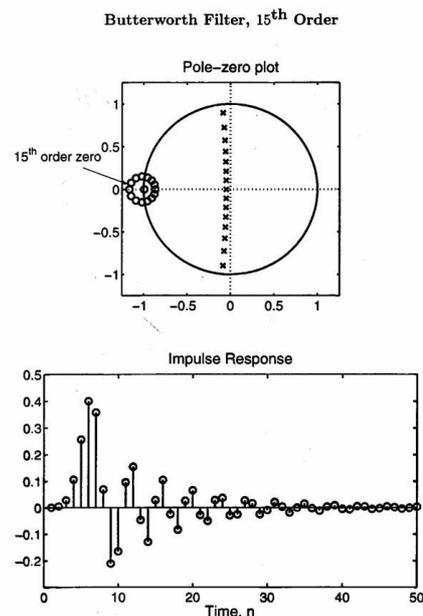


Figure 4: Pole-Zero Plot, Impulse Response

- Design Specifications:
1. Passband edge = 0.5π .
 2. Stopband edge = 0.6π .
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB.
 5. Maximum stopband gain = -30 dB.

- Direct form implementation of a Butterworth filter may cause instability because of coefficient quantization. The poles may move to the outside of the unit circle, and the zeros may spread around $z = -1$ as shown in the next figure. A solution is to use a cascade structure of first and second order systems, with zeros and poles grouped into complex conjugate pairs. Furthermore, since all of the zeros are located at $z = -1$, the numerator of the transfer function is $(1 + z^{-1})^N$. No multiplication is actually needed to implement the zeros, because this is equivalent to the cascade of N “delay and add” operations.



Design Specifications:

1. Passband edge = 0.5π .
2. Stopband edge = 0.6π .
3. Maximum passband gain = 0 dB.
4. Minimum passband gain = -0.3 dB.
5. Maximum stopband gain = -30 dB.

Chebyshev Filters

Type I

$$|H_{CH}(e^{j\omega})|^2 = \frac{1}{1 + \epsilon^2 V_N^2 \left(\frac{\tan(\omega/2)}{\tan(\omega_c/2)} \right)}$$

- Three design parameters: order N , PB cut-off frequency ω_c , allowed PB ripple ϵ (ie. the maximum allowable passband gain is 1, and the minimum allowable passband gain is $(1 - \epsilon)$).
- $V_N(x) = \cos(N \cos^{-1} x)$ is the N th-order chebyshev polynomial. Here we take the convention that $\cos^{-1} x$ becomes the inverse hyperbolic cosine and is imaginary if $|x| > 1$.

Some examples of lower order Chebyshev polynomials are:

$$\begin{aligned} V_0(x) &= 1 & V_1(x) &= x \\ V_2(x) &= \cos(2 \cos^{-1} x) = 2 \cos^2(\cos^{-1} x) - 1 = 2x^2 - 1 \end{aligned}$$

See Appendix B.2 of OSB for recurrence formulas for deriving Chebyshev polynomials.

- Equiripple in the passband, but decreases monotonically in the stopband.
- Similar to Butterworth filters, all zeros of a Chebyshev type I filter are located at $z = -1$. Following is a design example:

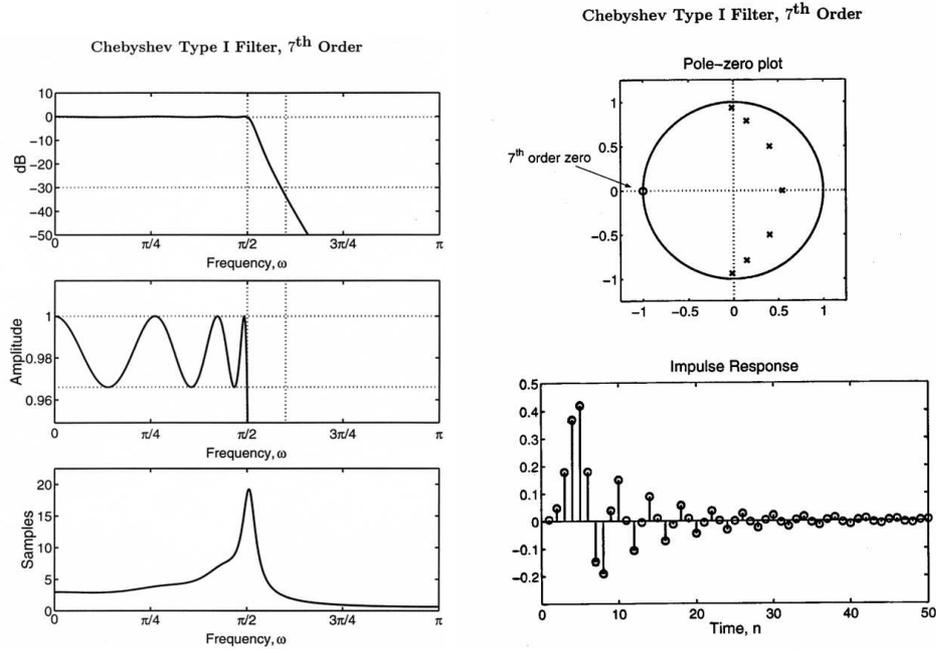


Figure 7: Magnitude Plot, Passband Detail, Group Delay

Design Specifications:

1. Passband edge = 0.5π .
2. Stopband edge = 0.6π .
3. Maximum passband gain = 0 dB.
4. Minimum passband gain = -0.3 dB.
5. Maximum stopband gain = -30 dB.

Figure 8: Pole-Zero Plot, Impulse Response

Design Specifications:

1. Passband edge = 0.5π .
2. Stopband edge = 0.6π .
3. Maximum passband gain = 0 dB.
4. Minimum passband gain = -0.3 dB.
5. Maximum stopband gain = -30 dB.

Type II

$$|H_{CH}(e^{j\omega})|^2 = \frac{1}{1 + \left[\epsilon^2 V_N^2 \left(\frac{\tan(\omega/2)}{\tan(\omega_c/2)} \right) \right]^{-1}}$$

- Monotonic in the passband; equiripple in the stopband.
- For a given set of specifications, the Chebyshev type I and Chebyshev type II design methods yield the same order.

- As shown in the figures below, poles are distributed around $z = 1$, similar to Butterworth designs, while zeros are arrayed on the unit circle at the stopband frequencies. Because zeros are not all located at $z = -1$, more multiplications are required in comparison to Chebyshev type I designs. OSB Example 7.5 compares Chebyshev type I and II filters in more detail.

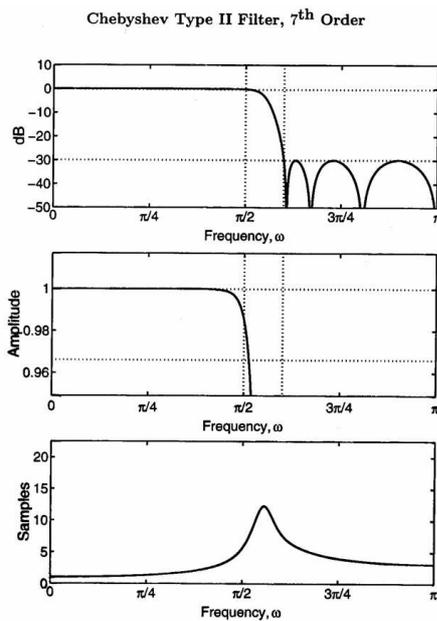


Figure 9: Magnitude Plot, Passband Detail, Group Delay

- Design Specifications:
1. Passband edge = 0.5π .
 2. Stopband edge = 0.6π .
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB.
 5. Maximum stopband gain = -30 dB.

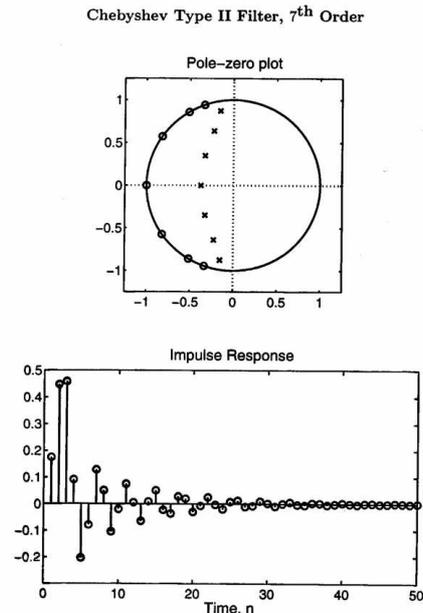


Figure 10: Pole-Zero Plot, Impulse Response

- Design Specifications:
1. Passband edge = 0.5π .
 2. Stopband edge = 0.6π .
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB.
 5. Maximum stopband gain = -30 dB.

- In both of the Chebyshev design methods, having a monotonic behavior in either the passband or the stopband suggests a lower order system might exist such that it satisfies the given set of specifications, but varies with equal ripple in both the passband and the stopband.

Elliptic Filters

- Four degrees of freedom: PB ripple, SB ripple, order, passband edge.
- Order of the system controls the transition bandwidth.
- Equiripple in both the passband and the stopband.
- Elliptic filters are the lowest order rational function approximation to a given set of magnitude specifications. All IIR filter designs we have discussed so far give nonlinear

phases, thus non-constant group delays. The greatest deviation from constant group delay occurs in all cases at either the edge of the passband, or within the transition band. In general, the Chebyshev type II approximation gives the smallest delay in the passband, and the widest region of the passband over which the delay is approximately constant. However, if constant group delay is not required, the elliptic approximation gives the lowest order system function.

- Following are some sample elliptic filter designs. Also see OSB Example 7.6.

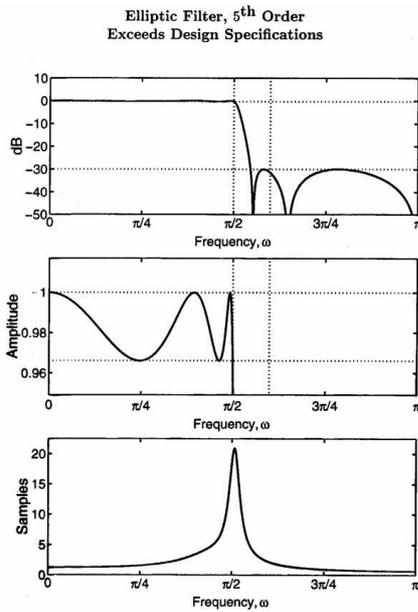


Figure 11: Magnitude Plot, Passband Detail, Group Delay

Design Specifications:

1. Passband edge = 0.5π .
2. Stopband edge = 0.6π .
3. Maximum passband gain = 0 dB.
4. Minimum passband gain = -0.3 dB.
5. Maximum stopband gain = -30 dB.

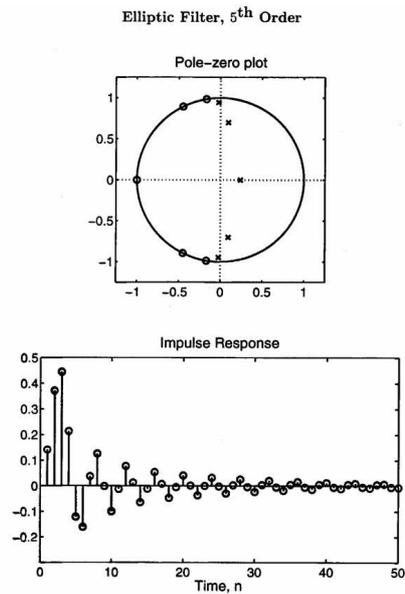


Figure 12: Pole-Zero Plot, Impulse Response

Design Specifications:

1. Passband edge = 0.5π .
2. Stopband edge = 0.6π .
3. Maximum passband gain = 0 dB.
4. Minimum passband gain = -0.3 dB.
5. Maximum stopband gain = -30 dB.

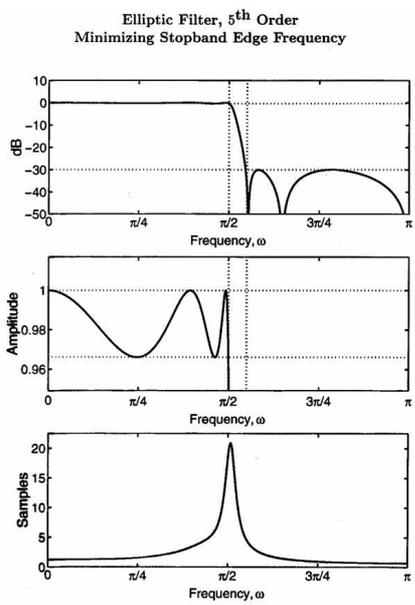


Figure 13: Magnitude Plot, Passband Detail, Group Delay

- Specifications Met:
1. Passband edge = 0.5π .
 2. **Stopband edge = 0.55π .**
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB.
 5. Maximum stopband gain = -30 dB.

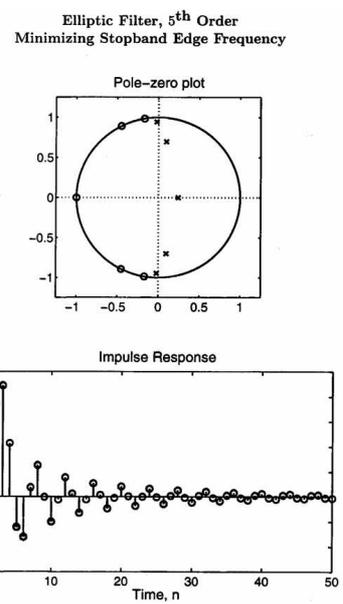


Figure 14: Pole-Zero Plot, Impulse Response

- Specifications Met:
1. Passband edge = 0.5π .
 2. **Stopband edge = 0.55π .**
 3. Maximum passband gain = 0 dB.
 4. Minimum passband gain = -0.3 dB.
 5. Maximum stopband gain = -30 dB.

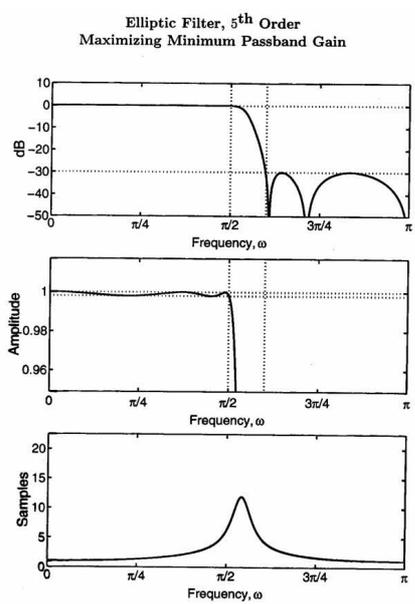
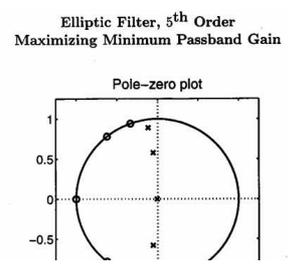


Figure 15: Magnitude Plot, Passband Detail, Group Delay

- Specifications Met:
1. Passband edge = 0.5π .
 2. Stopband edge = 0.6π .
 3. Maximum passband gain = 0 dB.
 4. **Minimum passband gain = -0.017 dB.**
 5. Maximum stopband gain = -30 dB.



Elliptic Filter, 5th Order
Minimizing Maximum Stopband Gain

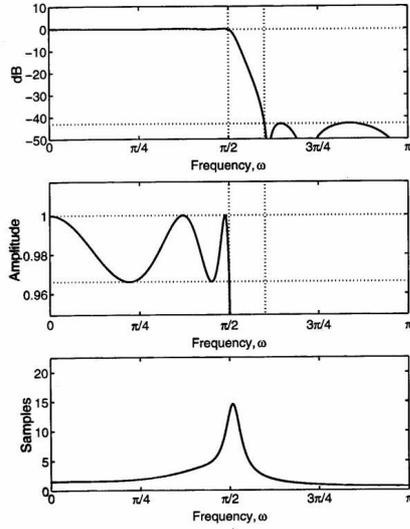


Figure 17: Magnitude Plot, Passband Detail, Group Delay

Specifications Met:

1. Passband edge = 0.5π .
2. Stopband edge = 0.6π .
3. Maximum passband gain = 0 dB.
4. Minimum passband gain = -0.017 dB.
5. Maximum stopband gain = -43 dB.

Elliptic Filter, 5th Order
Minimizing Maximum Stopband Gain

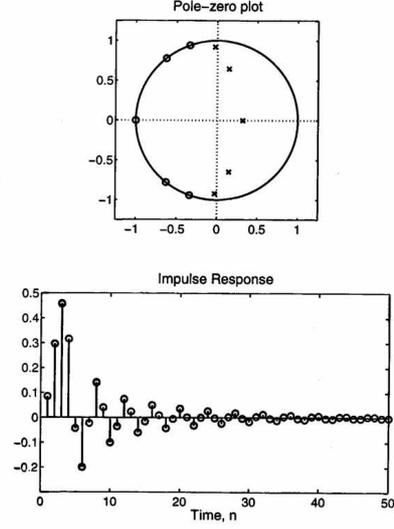


Figure 18: Pole-Zero Plot, Impulse Response

Specifications Met:

1. Passband edge = 0.5π .
2. Stopband edge = 0.6π .
3. Maximum passband gain = 0 dB.
4. Minimum passband gain = -0.017 dB.
5. Maximum stopband gain = -43 dB.