Department of Engineering

Lecture 09: Filters

Matthew Spencer

Harvey Mudd College

E157 – Radio Frequency Circuit Design

1

Department of Engineering

Filter Specifications and the Filter Prototype Function

Matthew Spencer

Harvey Mudd College

E157 – Radio Frequency Circuit Design

2

In this video we're going to start talking about filters by defining a language that we use to describe them.

Department of Engineering

Filters are Like Extended Matching Networks

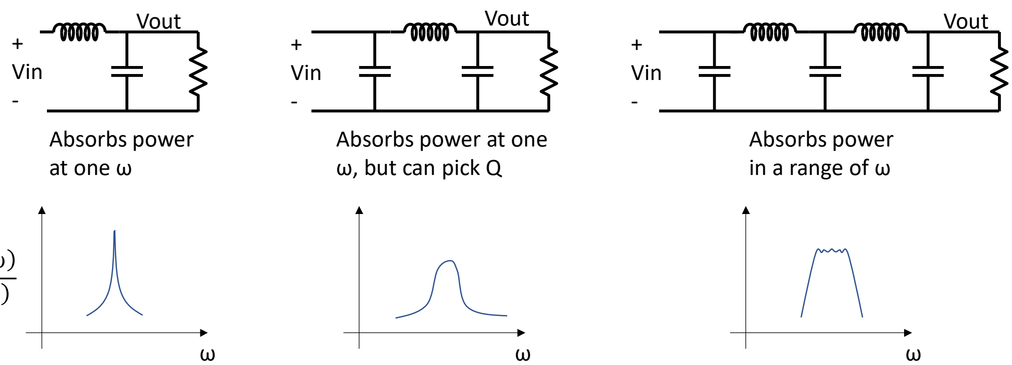

$$H(j\omega) = \frac{V_{out}(j\omega)}{V_{in}(j\omega)}$$

Filters are a natural follow on after talking about matching networks because you can think of them as an extension of the same idea. We showed that an L-match lets us absorb energy at one frequency (which is resonance) and reflect it at every other frequency, so we could think of an L match as a type of filter. That's particularly obvious if we define a transfer function across an L match network from Vin to Vout, which would look like a narrow resonant peak. Adding more components in a pi match allowed us to control the shape of that peak and smear it out over more frequencies. So it stands to reason that by adding even more components to our matching network, we could control whether a signal is passed or reflected over a wider frequency. That turns out to be true, and the type of circuit that achieves this frequency response is referred to as an LC ladder filter. We opt to make LC ladder filters out of inductors and capacitors for the same reason we use inductors and capacitors in matching networks: lossless components make the filter lossless.

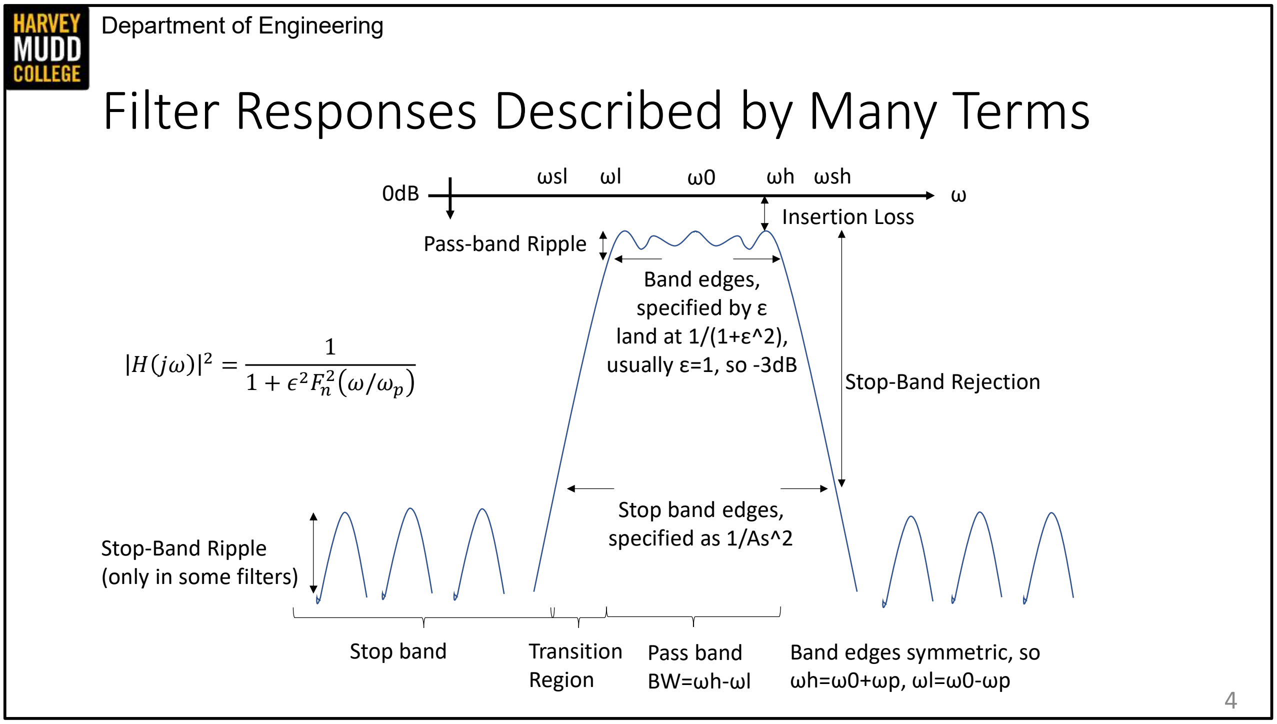

The exact shape of the filter transfer function can be captured by this equation, which we call the filter prototype because it is highly customizable. The epsilon parameter affords some flexibility, but the real special sauce is the Fn function, which represents a ratio of nth order polynomials that varies between zero and one for a while, then flies off to infinity. The order of this polynomial affects the steepness of the filter behavior and the complexity of the filter design, so it's important. The order of that polynomial is also referred to as the order of the filter.

CLICK this is a general picture of what this transfer function might look like. You can see that the x-axis indicates that filter behavior changes with frequency, and that the y-axis is has units of dB, which means that we're framing our discussion in terms of how much power gets through a filter. That's consistent with the picture painted by our equation, which represents a transfer function squared, and accordingly describes a ratio of powers rather than a ratio of voltages or currents.

CLICK the filter is described by a center frequency w0, which sits in the middle of a frequency region called the pass band. The upper and lower edges of this band are described by frequencies wh and wl, and they are symmetric about w0, so we can specify wh=w0+wp and wl=w0-wp.

CLICK There is some loss in the pass band, which is referred to as insertion loss. The transfer function also varies in the pass band, and that behavior is referred to as pass-band ripple. The pass band ripple measurement is the difference between the insertion loss and the lowest allowed value of the pass band. Some filters have pass bands that only ever decrease, which means we can use the vocabulary word monotonic to describe them. Those filters don't have ripple in the sense you might think of it, but we still call the steady decay of the transfer function ripple.

CLICK In a properly designed filter, the maximum level of pass-band ripple will occur at the edges of the pass band, wh and wl. You can pick exactly how much ripple you want to allow in the pass band by setting the epsilon parameter because the Fn function stays between +/-1 in the pass band. That means the band edges are defined by the point at which Fn=1, and the transfer function is 1/(1+e^2) at the band edge. Usually epsilon is set to 1 so that the squared transfer function has a value of 1/2 at the band edge, signaling that the power has been cut in half.

CLICK You are free to pick how much rejection you want outside of the pass band, and that quantity is called the stop-band rejection. You usually specify this with a quantity As, so a stop band rejection of 20dB would correspond to an As of 10. Stop band rejection is measured relative to the top of the pass band.

CLICK The frequencies at which the transfer function reaches the stop band rejection are called the stop-band edges, and we give them the symbols wsh and wsl. The frequencies between the pass band edge and the stop band edge are called the transition region.

CLICK Finally, the frequency band outside of the transition region is called the stop band. Some filters still have finite transfer function values in this region, and that behavior is referred to as stop-band ripple. The stop band ripple sets a maximum value for stop-band rejection.

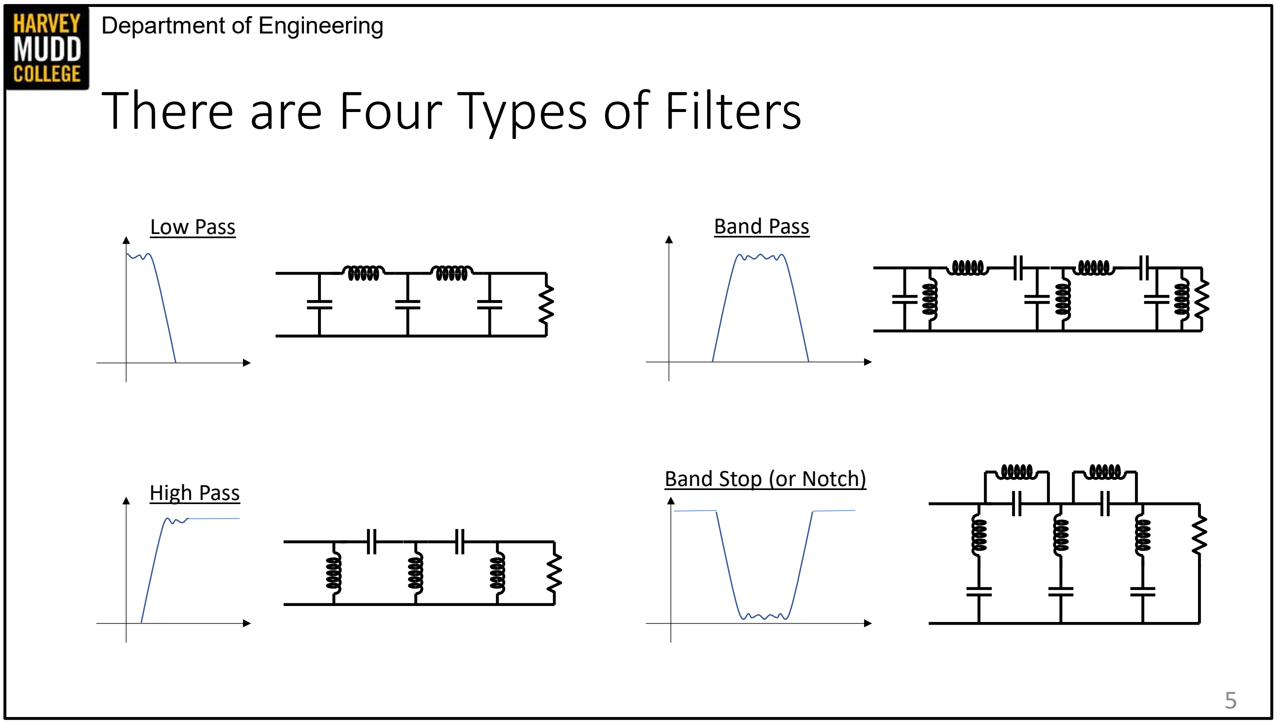

The picture we've been looking at thus far is a band pass filter, but the filter prototype can be used to describe all four common types of filters, which, as a reminder, are low-pass, high-pass, band-pass and band-stop. You can build lossless LC ladder versions of each of these filters, and I've drawn them on this slide.

Thinking about the asymptotic impedance of each component can give you some intuition about how the low-pass and high-pass filters I've drawn exhibit the low-pass and high-pass behavior. Consider the low pass filter as a starting point: at low frequencies the capacitors are open and the inductors are shorts, allowing input voltages to reach the load easily, and at high frequencies the capacitors are shorts and the inductors are open, which zeroes the voltage on the load.

Intuition for the band-pass and band-stop circuits requires thinking about the behavior of their LC tanks at resonance. Parallel tanks ring open and series tanks ring short, so the band bass will have series short circuits and shunt open-circuits at resonance, which allows voltages to reach the load. Off resonance, either the inductor or the capacitor will block signals in the series elements and absorb signals in the shunt elements.

It's worth noting that all the filters on this slide are 3rd order, they would produce 3rd order polynomials in the filter prototype transfer function. This is related to the fact that the low-pass and high-pass filters have three independent energy storage elements, which are the three shunt elements. The series elements aren't independent of the shunt elements because they are part of LC tanks.

Summary

- Filters absorb power over a frequency band and reflect it elsewhere.
- Filters are described by a prototype w/ a polynomial and an order

$$|H(j\omega)|^2 = \frac{1}{1 + \epsilon^2 F_n^2(\omega/\omega_p)}$$

- Filter have many parameters: bandwidth, stop-band rejection, etc.
- There are four types of filters, they are made by rearranging L and C

Department of Engineering

Laplace Interpretation of Filter Design, Filter Prototype

Matthew Spencer
Harvey Mudd College
E157 – Radio Frequency Circuit Design

7

In this video we're going to look at a graphical method of calculating the magnitude of transfer functions that sheds some light onto how filters work. Filter design can be a little opaque because the circuit theory requires calculating the values of loaded components that give rise to weird polynomials in transfer function, but I find this graphical interpretation makes filters simpler to understand.

You Can Calculate |H(jw)| Graphically

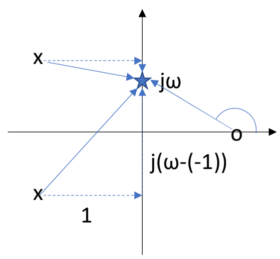

$$H(s) = \frac{1}{s^2 + 2s + 2} = \frac{1}{(s+1+j)(s+1-j)}$$

$$H(j\omega) = \frac{1}{(j\omega+1+j)(j\omega+1-j)} = \frac{1}{\big(1+j(\omega-1)\big)\big(1+j(\omega+1)\big)}$$

$$|H(j\omega)| = \frac{1}{\sqrt{1+(\omega-1)^2}\sqrt{1+(\omega+1)^2}}$$

$$|H(j\omega)| = \frac{\prod_{zeros}|j\omega - s_{zero}|}{\prod_{poles}|j\omega - s_{pole}|}$$

Aside on RHP zeros: $\angle H(j\omega) = \sum_{zeros} \angle(j\omega - s_{zero}) - \sum_{poles} \angle(j\omega - s_{pole})$

8

The graphical method for calculating filter behavior has its roots in pole-zero plots. So we're starting to understand this method by drawing a pole-zero plot for this transfer function. I've taken the liberty of doing the quadratic equation for you on this first line, which reveals that we have poles at -1-j and -1+j, and I plotted those on the left. As a note, this is an example of a second order filter because there are two poles on the pole-zero plot. We're only looking at H(s) and not |H(s)|^2, so we can't say for sure what Fn is yet, but it's probably also quadratic.

CLICK If we want to find the value of this transfer function at a point jw, then we can just substitute jw for s in our transfer function. This is equivalent to evaluating the transfer function at some point on the imaginary axis because jw is just a locus of points on the s plane. That's cool enough that it bears repeating: the frequency response function for this transfer function is just the values along the imaginary axis, and we can find the frequency response function by evaluating the transfer function at every point on the imaginary axis.

CLICK We can group like terms in the denominator of this expression, which starts to have a geometric interpretation. Each of these regrouped poles specifies a vector that points from a pole to the point jw. For instance, the bottom pole is a distance 1 from the jw point in the real direction, and a distance 1+w from the pole in the imaginary direction.

CLICK We can take the magnitude of this expression, and in doing so we find that the denominator of this magnitude is the product of the lengths of the vectors from the poles to jw. CLICK That generalizes to any transfer function: the magnitude at a point on the s-plane, which I've labeled as jw in our case, is the product of the lengths of the vectors that point from zeros to the point in the question divided by the product of the lengths of the vectors pointing from poles to the point in the question. So we can understand the behavior of our magnitude just by looking at the lengths of a bunch of vectors on the s-plane.

CLICK We're going to apply that result to filters on the next slide, but I need to put in a sidebar about calculating phase graphically. There's a similar graphical technique for finding phase, where the phase of the transfer function is the sum of the angles from zeros to the point in question minus the sum of angles from poles to the point in question. That's a cool thing to remember, and it explains the baffling behavior of right half plane zeros. Right half plane zeros contribute positive magnitude and negative phase to Bode plots, which is weird, and also really bad for stability in feedback. You can see that behavior on this plot, because increasing omega will make the angle from the right half plane pole smaller, shrinking the the total phase of the transfer function.

HARVEY MUDD COLLEGE
Department of Engineering

Filters Have Many Poles at Same Frequency

$$|H(j\omega)| = \frac{\prod_{zeros}|j\omega - s_{zero}|}{\prod_{poles}|j\omega - s_{pole}|}$$

$$|H(j\omega)|^2 = \frac{1}{1+\epsilon^2 F_n^2(\omega/\omega_p)}$$

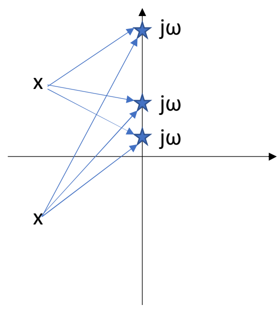

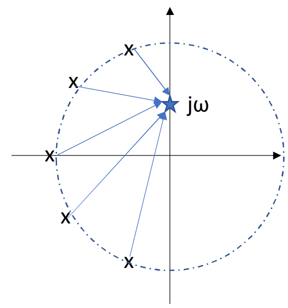

$F_2(\omega/\omega_p) = (\omega/\omega_p)^2$

$F_5(\omega/\omega_p) = (\omega/\omega_p)^5$

9

We can apply this graphical technique to understand how filters work. I've copied over the second order filter from the previous page and drawn a point omega in between the two poles.

CLICK When I look at a different input frequency between the poles, we can see that the distance to one pole gets larger while the distance to the other pole gets shorter. That results in the overall transfer function staying flat.
CLICK However, if I look at a value of j*omega outside of the two poles, we find that both of the vectors to the poles are getting longer as omega increases. That means the value of the transfer function will start falling off quickly.
CLICK And if I wanted to take advantage of that behavior, I could add a lot more poles at about the same distance from the origin. In this example of a fifth order filter, five poles have vectors that cancel out for omega points inside a circle in the s-plane, and if the j*omega point moved outside the circle, we'd have five vectors getting longer and decreasing the transfer function.
CLICK This change, where we add more poles to cause transfer functions to fall off more quickly, corresponds to increasing the order of the filter, and I've included the Fn polynomials for these filters to show that. The second order filter uses the F2 polynomial, while the fifth order filter uses the F5 polynomial.

This is a low-pass filter, which we've seen because the transfer function gets small when omega gets big. We'd have to move poles and zeros around in a dramatic way to make a high-pass filter – in fact we usually do that by inverting omega and omega p in the prototype function, which is a non-linear transform. We'd need a different dramatic change for a bandpass or band stop filter: either of those would require twice as many poles.

Department of Engineering

Different Filters Space Poles/Zeros Differently

$$|H(j\omega)| = \frac{\prod_{zeros}|j\omega - s_{zero}|}{\prod_{poles}|j\omega - s_{pole}|}$$

$$|H(j\omega)|^2 = \frac{1}{1 + \epsilon^2 F_n^2(\omega/\omega_p)}$$

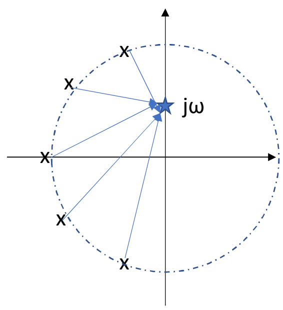

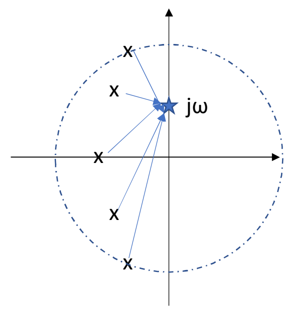

$$F_n(\omega/\omega_p) = (\omega/\omega_p)^n$$

$$F_n(\omega/\omega_p) = C_n(\omega/\omega_p)$$

10

Last slide we saw how changing the filter order changed the behavior of filters. We can also change the behavior of filters by changing the location of the poles. In the filter on the right, the poles are located on an ellipse instead of circle, which will result in the length of the vectors to the omega point cancelling imperfectly in the pass band, but it also means that vectors get longer faster outside of the pass band, so we've traded some pass band ripple for a faster rolloff in this arranagement of poles, even for filters with the same order. Positioning poles in this way required us to change the prototype polynomial to something called an nth order Chebyshev polynomial, which we'll talk more about soon.

Summary

- Frequency responses can be calculated graphically as a product of magnitudes of distance from poles to points $j\omega$ on the imaginary axis.
- Filters have many poles at similar frequencies to achieve quick rolloff.
- The arrangement of those poles is determined by the filter prototype polynomial.

Department of Engineering

Filter Examples

Matthew Spencer
Harvey Mudd College
E157 – Radio Frequency Circuit Design

12

In this video we're going to look at a few common filters and we'll compare their frequency responses, pole-zero plots and filter prototype functions. I'd like you to pay special attention to features of these filters that will let you pick them out on sight.

HARVEY MUDD COLLEGE
Department of Engineering

Butterworth Poles Are Roots of Unity

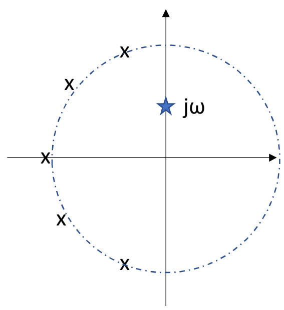

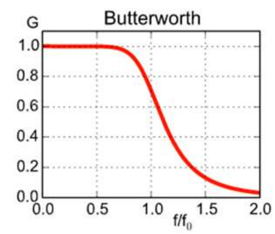

$$|H(j\omega)|^2 = \frac{1}{1 + \epsilon^2 F_n^2(\omega/\omega_p)}$$

$$F_n(\omega/\omega_p) = (\omega/\omega_p)^n$$

Nth roots of unity (on LHP)

- Flat pass band, then rolls off at -20n dB/decade.
- “maximally flat”, Though Chebyshev II filters technically flatter.

The Butterworth filter is one of the most common filters, and one of the easiest to understand. Butterworth filters have prototype polynomials of omega over omega_p raised to the nth order, which means that the poles are going to be given by half of the 2n roots of unity on the omega p circle. We know that because the roots of the filter prototype’s denominator are given by 1=(w/wp)^2n . The filters are represented by the half of the roots of unity on the left half plane, which is perhaps unsurprising because the right half plane roots of unity would indicate an unstable system.

The roll off of the butterworth filter is just the combined effect of a bunch of co-located poles, so the transfer function rolls off at order * 20 dB/decade.

Butterworth filters are often called maximally flat because they have monotonically decreasing ripple in the pass band. However, type II Chebyshev filters are technically a little flatter in the pass band for any given value of epsilon.

Department of Engineering

Chebyshev I Poles Are On an Ellipse

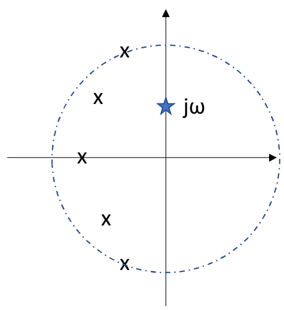

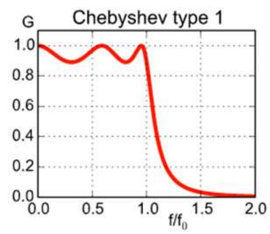

$$|H(j\omega)|^2 = \frac{1}{1+\epsilon^2 F_n^2(\omega/\omega_p)}$$

$$F_n(\omega/\omega_p) = C_n(\omega/\omega_p)$$

Chebyshev Polynomials.

$$C_n(x) = 2xC_{n-1}(x) - C_{n-2}(x)$$

- Ellipse causes faster roll off at cost of ripple in pass band.

14

Type I Chebyshev filters have poles that are placed on an ellipse, which results in ripple in the pass band, but contributes to a sharper rolloff than Butterworth filters. The prototype polynomial for a Chebyshev filter is a nth order Chebyshev polynomial, which is recursively defined as shown on the slide. That definition fails for the first two polynomials, but, for reference, C0=1 and C1=x.

HARVEY MUDD COLLEGE
Department of Engineering

Chebyshev II Adds Zeros to the Filter

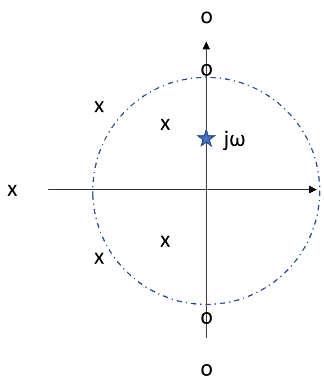

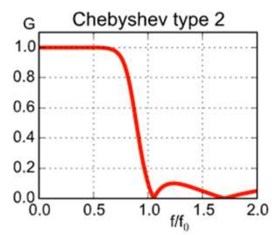

$$|H(j\omega)|^2 = 1 - \frac{1}{1+\epsilon^2 F_n^2(\omega/\omega_p)}$$

$$F_n(\omega/\omega_p) = C_n(\omega_p/\omega)$$

Upside down Chebyshev Polynomials in numerator & denominator

- Synthesized by taking 1-Cn(ωp/ω) to swap pass and stop band
- Zeros provide nulls for fast roll off, but create stop band ripple.

Type II Chebyshev filters are notable because they let us introduce two new ideas: low-to-high pass transformations and nulling zeros. The motivation for Type II Chebyshev filters is to swap the ripple in the pass band of the type I filter for ripple in the stop band. You can do that by taking 1 minus the Type I Chebyshev transfer function, but that gives you a high pass filter. However, if you make your Chebyshev polynomials a function of omega p over omega, you swap high and low frequencies in the filter prototype function again. This results in the Type II Chebyshev, which has poles placed according to the inverse of an ellipse, and also has zeros on the y axis.

The zeros on the y axis are interesting, because our transfer function has to be zero when our omega point is on one of those zeros. That means you can use them to cause the value of the function to collapse very quickly, and the zeros are responsible for the sharp rolloff of this filter. However, zeros also cancel out poles in the transfer function at some frequencies, which results in a finite stop band instead of one that rolls off forever.

HARVEY MUDD COLLEGE Department of Engineering

Elliptic Filters Have Ripple in Pass & Stop Band

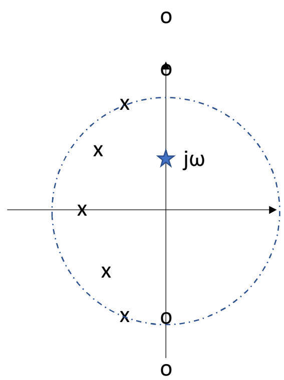

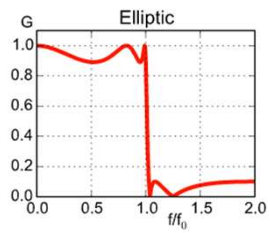

$$|H(j\omega)|^2 = \frac{1}{1+\epsilon^2 F_n^2(\omega/\omega_p)}$$

$$F_n(\omega/\omega_p) = R_n(\omega/\omega_p, \xi)$$

Elliptic rational function with selectivity ξ

- Ripple in pass band and stop band → very sharp rolloff
- Reduces to Chebyshev I, Chebyshev II, or Butterworth as ripple drops

Elliptic filters, which are also called Cauer filters sometimes, combine the features of type I and type II Chebyshev filters to achieve very sharp rolloff in the transition region. They have ripples in the pass band from poles placed on an ellipse and they also have nulling zeros. These features are specified by their prototype function, which is called an elliptical rational function. The elliptical rational function has a selectivity factor which we label xi, and tuning that selectivity factor can cause the behavior of elliptic function to reduce to a Chebyshev I filter, a Chebyshev II filter or a Butterworth filter.

Harvey Mudd College Department of Engineering

Other Filters

- Bessel(-Thomson) Filters have flat phase responses, constant delay.
- Can make filters out of transmission lines of different widths called constant-k filters.
- Q minimizing filter variants reduce the effect of component variation

17

We've only done a quick tour of filters here and there are other types in the world. Bessel filters have flat phase responses that give rise to constant group delay. Constant-k filters can be built out of transmission line segments, and Q-minimizing filters try to reduce the quality factor of the resonators comprising a filter in order to make it more resistant to randomly varying components.

Summary

- There are many kinds of filters that can be identified by steepness of roll off and presence of ripple in pass/stop bands
  - Butterworth: no ripple
  - Chebyshev I: pass band ripple
  - Chebyshev II: stop band ripple
  - Elliptic: ripple in both pass and stop bands

- Zeros null functions quickly, but result in stop band ripple

- Filter shapes are linked to the filter prototype polynomials

Department of Engineering

Filter Design Process and Filter Tables

Matthew Spencer

Harvey Mudd College

E157 – Radio Frequency Circuit Design

19

In this video we're going to talk about designing a filter, which is going to require us to use a calculation tool called a filter table.

HARVEY MUDD COLLEGE
Department of Engineering

Pass & Stop Ripple, ε, As, ωp→ Type, Order

- Get epsilon from allowable passband deviation: $\epsilon = \sqrt{10^{0.1\delta} - 1}$
- Convert As and ε to filter order from equation or plot
  - Butterworth → $n = \ln(A_s/\epsilon)/\ln(\omega/\omega_p)$
  - Chebyshev → $n = \cosh^{-1}(A_s/\epsilon)/\ln(\omega/\omega_p)$
  - Elliptic → elliptic integrals
- Can double check by looking up or simulating filters (lots of resources)

20

The first step in design is to pick the features we want in our filter, which includes the size of the passband ripple, the stopband rejection, and the bandwidth. We'll put these features through a mathematical transformation to find the filter order.

Epsilon tells you the total allowable ripple in the pass band, but it's a linear representation and pass band ripple is often given by a deviation in dB. You can convert between the two by converting the dB representation back into a linear scale as shown on this slide, where delta is the amount of ripple in dB. This comes from setting 1/(1+e^2) equal to the maximum ripple.

Once you have epsilon and you've picked your stop-band rejection based on your application, you can combine those numbers to pick your filter order. The calcaulations I show here are approximations based on filters rolling off at order times 20dB per decade. I didn't show an equation for elliptic filters because finding the order of an elliptic filter requires something called an elliptic integral.

You can check your work in picking your filter properties with a variety of simulators and websites. Lots of people have implemented filter calculators.

Department of Engineering

Use a Filter Table to Synthesize a Low-Pass

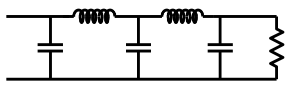

OR

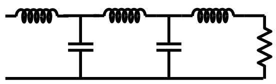

Order	R_S	C_1 a_1	L_2 a_2	C_3 a_3	L_4 a_4	C_5 a_5	L_6 a_6	C_7 a_7
1	1.0	2.0000						
2	1.0	1.4142	1.4142					
3	1.0	1.0000	2.0000	1.0000				
4	1.0	0.7654	1.8478	1.8478	0.7654			
5	1.0	0.6180	1.6180	2.0000	1.6180	0.6180		
6	1.0	0.5176	1.4142	1.9319	1.9319	1.4142	0.5176	
7	1.0	0.4450	1.2470	1.8019	2.0000	1.8019	1.2470	0.4450

$$Z(\omega)/Z_0 = Z_{normalized}(1\ \text{rad/s}) \begin{cases} \omega_p L = Z_0 (1\ \text{rad/s}) L_{normalized} \\ \dfrac{1}{\omega_p C} = \dfrac{Z_0}{(1\ \text{rad/s}) C_{normalized}} \end{cases}$$

- Need n energy storage elements to create n poles
- L and C values have already been found, codified in filter tables
- Filter tables are normalized to Zin=1 ohm and ωp=1 rad/s

21

Now that we have the order, we can use that to pick component values such that poles wind up in the right spot. There's a bunch of interesting polynomial math that goes into calculating L and C values, but it's all been done before and we're going to take advantage of that. The results of these filter calculations are tabulated in resources called filter tables, which give you values for inductors and capacitors that will synthesize a low-pass filter of your desired order, with an input impedance of 1 ohm and a bandwidth of 1 rad/s. Since we usually want different input impedances and bandwidths, we need to "un-normalize" the values in the filter table. You do that using the formulas shown here, which I remember by thinking about impedance. That is to say, the impedance I want at the bandwidth I want divided by Z0 has to equal the normalized impedance, has to equal the normalized impedance evaluated at 1 rad/s.

Also, as a note, there are two forms of low pass filter that you could choose to synthesize, an inductor first version and a capacitor first version. Both are valid low pass filters and they actually use the same filter table values. I've drawm examples of the two types of 3rd order filter you can make. Also, for reference, this filter table is for Butterworth filters.

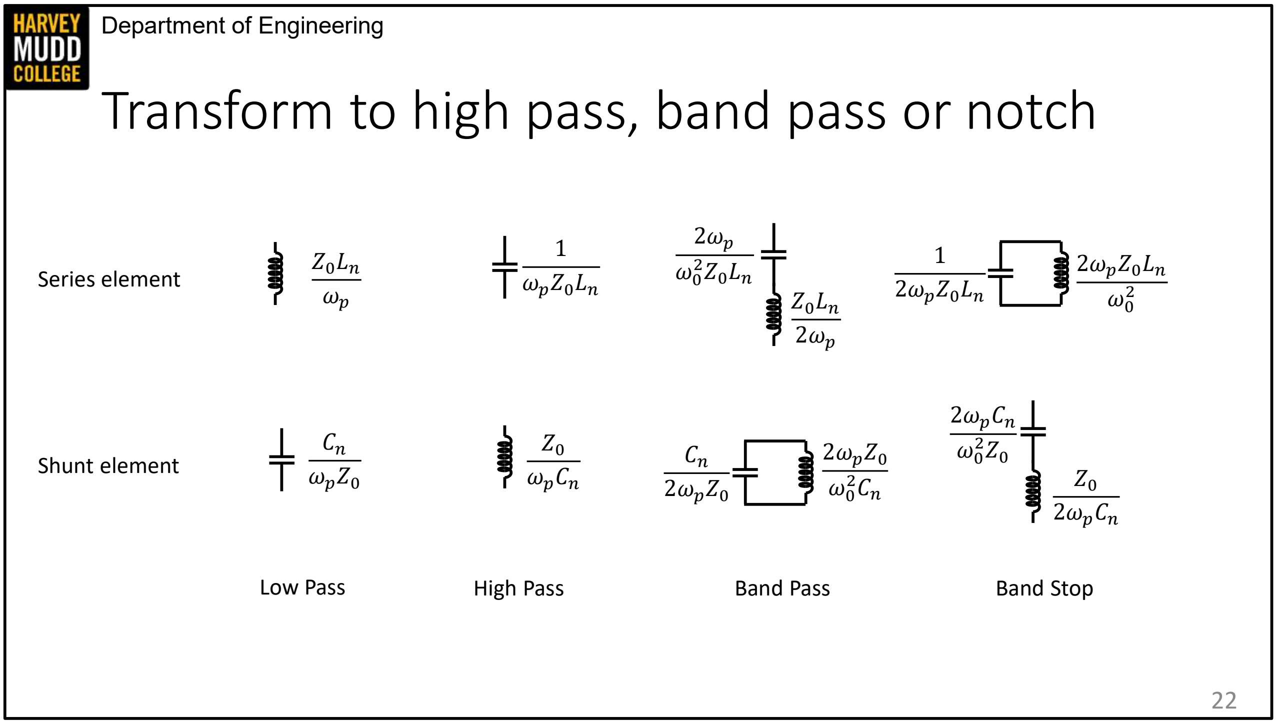

If you want to make a high-pass, band-pass or band-stop filter, you do so by substituting new series and shunt elements into your low-pass filter. This slide shows the components that you would use as series and shunt elements in each filter type, along with formulas for un-normalizing them. Here the subscript n indicates the normalized impedance from the filter table. Note that 2wp is the bandwidth of bandpass or band stop filters, and that w0 is the center frequency.

Summary

- Find filter type and order from ripple, ε, As, ωp
- Make a low pass filter from filter tables
- Un-normalize filter tables according to Z/Z0=Zn(1 rad/s)
- Transform from low-pass filter to other filter types with substitutions