Contents

Chapter 10 — Online Content

Articles

- 6-Meter Filter with Harmonic Suppression by Paul Wade, W1GHZ
- A High Performance, Low Cost 1.8 to 54 MHz Low Pass Filter by Bill Jones, K8CU
- Altoids Tin Filters by Paul Wade, W1GHZ
- Clean Up Your Signals with Band-Pass Filters — Parts 1 and 2 by Ed Wetherhold, W3NQN
- Combline Filters for VHF and UHF by Paul Wade, W1GHZ
- Crystal Filter Design and Crystal Characterization
- Crystal Parameter Measurements Simplified by Chuck Adams, K7QO
- Hands On Radio: *ELSIE* Filter Design — Parts 1 and 2 by Ward Silver, NØAX
- HF Yagi Triplexer Especially for Field Day by Gary Gordon, K6KV
- High-Power Harmonic Filters by George Cutsogeorge, W2VJN
- High-Power HF Band-Pass Filter Design by Jeff Crawford, KØZR
- Improved Audio-Frequency Bandpass Filter for Morse Code Reception by Jim Tonne, W4ENE
- Manual Filter Design Examples by Jim Tonne, W4ENE
- SDR Simplified — Cascaded Integrator Comb Filters by Ray Mack, W5IFS
- SDR Simplified — Filter Design Program by Ray Mack, W5IFS
- SDR Simplified — Fourier Transforms by Ray Mack, W5IFS
- SDR Simplified — More Filter Activities by Ray Mack, W5IFS
- Using Active Filter Design Tools by Dan Tayloe, N7VE

Design Software

The following Windows software by Jim Tonne, W4ENE, is available with the online content.

- *Diplexer* for design and analysis of diplexer filters
- *Elsie* for design and analysis of lumped-element LC filters
- *Helical* for design and analysis of helical-resonator bandpass filters
- *SVC Filter Designer* for design and analysis of lumped element high-pass and low-pass filters
- *QuadNet* for design and analysis of active all-pass networks for SSB operation

Chapter 10
Analog and Digital Filtering

This chapter discusses the most common types of filters used by radio amateurs, both analog and digital. Design information is supplied where appropriate or references to software or detailed design procedures are supplied.

The sections describing basic concepts, lumped element filters and some design examples were initially prepared by Jim Tonne, W4ENE, and updated by Ward Silver, NØAX. Digital filter material was updated by Doug Grant, K1DG, based on material originally developed by Alan Bloom, N1AL. Additional digital filter material was based on the SDR: Simplified column in *QEX* by Ray Mack, W5IFS. Material on cavity filters was provided by John Portune, W6NBC.

The online information's design example for active filters was provided by Dan Tayloe, N7VE. The online information's design example for crystal filters was developed by Dave Gordon-Smith, G3UUR.

10.1 Introduction

Electrical filters are circuits used to process signals based on their frequency. For example, most filters are used to pass signals of certain frequencies and reject others. The electronics industry has advanced to its current level in large part because of the successful use of filters. Filters are used in receivers so that the listener can hear only the desired signal; other signals are rejected. Filters are used in transmitters to pass only one signal and reject those that might interfere with other spectrum users. **Table 10.1** shows the usual signal bandwidths for several signal types.

The simplified receivers in **Figure 10.1** show filter use in both an analog, superheterodyne receiver (Figure 10.1A) and a simplified SDR receiver (Figure 10.1B). A *preselector* filter is placed between the antenna and the receiver's front end to pass all frequencies within a given amateur band with low loss. Strong out-of-band signals from broadcast, commercial, or military stations are rejected to prevent them from overloading the receiver input. The preselector filter is almost always built with *lumped-element* or "LC" technology.

In an analog receiver, there are one or more intermediate frequency (IF) filters to select only the desired signal. The IF filter closest to the antenna may be a relatively wide "roofing" filter with a bandwidth of several kHz to reject strong in-band signals not close to the desired signal frequency. Following detection or demodulation, an audio filter is placed somewhere ahead of the AF amplifier and speaker to rejects unwanted products from the noise. The audio filter is often implemented with *active filter* technology. For an SDR receiver, similar preselector and audio filters may be used but the filters which select a single signal are implemented as *digital filters* in the I/Q Processing block.

The complementary transmitter block diagrams are shown in **Figure 10.2** in which a similar array of filters appears in reverse order. In the analog transmitter, an audio filter between the microphone and the balanced mixer rejects noise and unwanted speech components. Since the balanced mixer generates both lower and upper sidebands, an IF filter is placed at the mixer output to pass only the desired lower (or upper) sideband. In the SDR version, the input signals are digitized and processed before the I/Q signals are filtered for transmission.

Finally, a filter at the output of the transmit mixer passes only signals within the amateur band in use rejects unwanted frequencies generated by the mixer to prevent them from being amplified and transmitted. Filters at the transmitter output attenuate the harmonics of the transmitted signals.

Table 10.1
Typical Filter Bandwidths for Typical Signals

Source	*Required Bandwidth*
Fast-scan analog television (ATV)	4.5 MHz
Broadcast-quality speech and music	15 kHz (from 20 Hz to 15 kHz)
Communications-quality speech	3 kHz (from 300 Hz to 3 kHz)
Slow-scan television (SSTV)	3 kHz (from 300 Hz to 3 kHz)
HF Digital (general)	100 to 1500 Hz (depends on modulation and bit rate)
HF RTTY (standard shift)	250 to 500 Hz
Radiotelegraphy (Morse code, CW)	200 to 500 Hz
PSK31 digital modulation	100 Hz

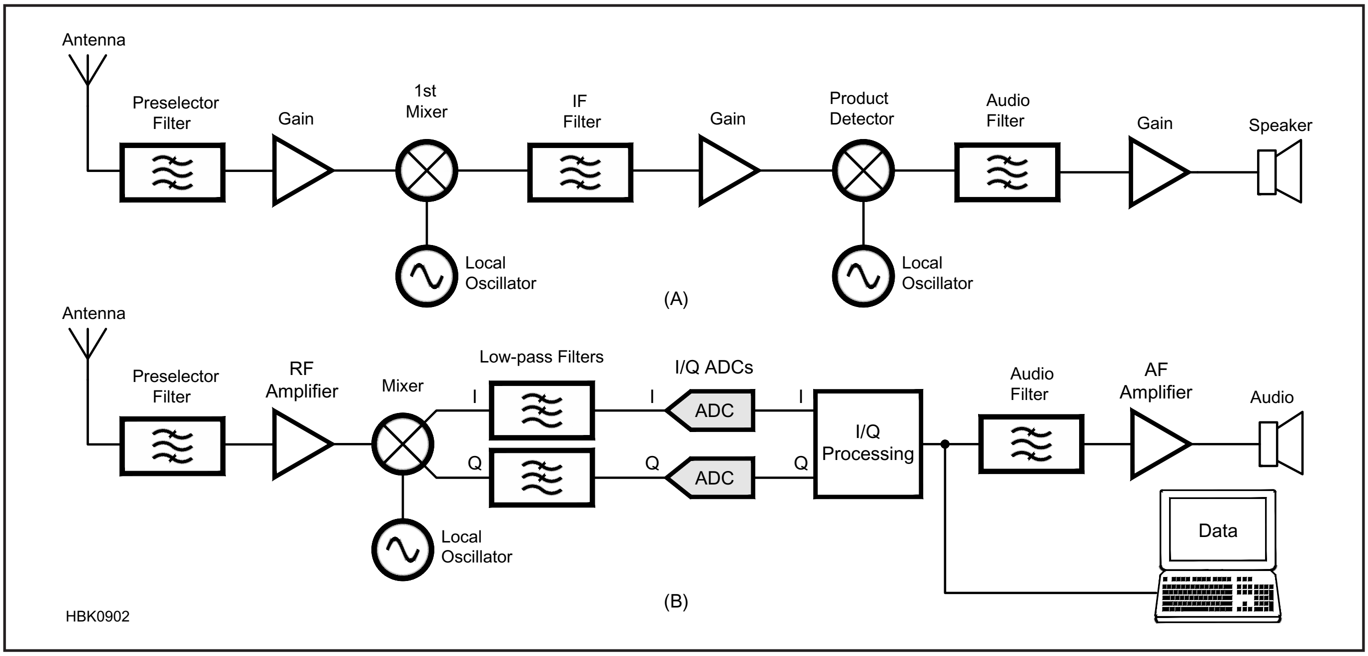

Figure 10.1 — Single-band SSB superheterodyne receiver (A) and SDR SSB/Data receiver (B) showing the filters typically used by each.

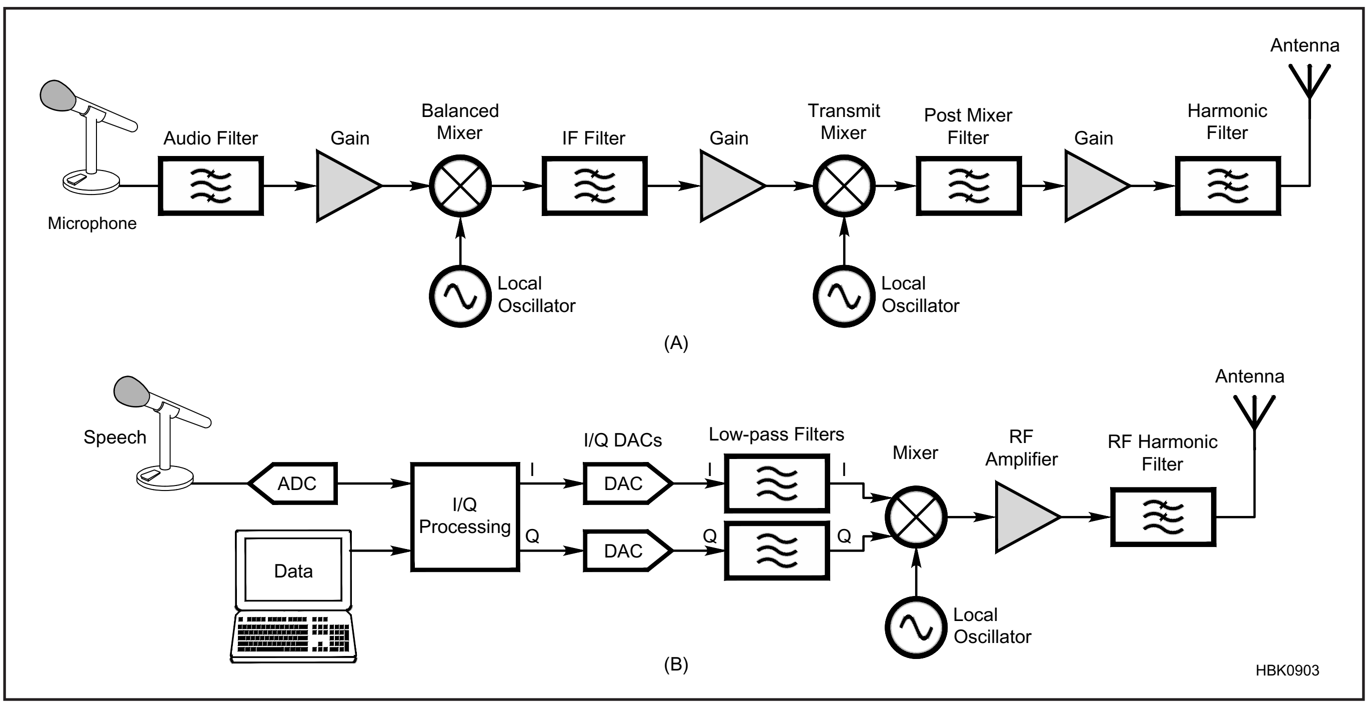

Figure 10.2 — Single-band SSB superheterodyne transmitter (A) and SDR SSB/Data transmitter (B) showing the filters typically used by each.

10.2 Filter Basics

10.2.1 Filter Magnitude Responses

A common type of filter, the *band-pass*, passes signals in a range of frequencies — the *passband* — while rejecting signals outside that range — the *stopband*. To pass signals from dc up to some *cutoff frequency* at which output power is halved (reduced by 3 dB) we would use a *low-pass* filter. To pass signals above a cutoff frequency (also called the "3-dB frequency") we would use a *high-pass* filter. Similarly, to pass signals within a range of frequencies we would use a *band-pass* filter. To pass signals at all frequencies *except* those within a specified range requires a *band-stop* or *notch filter*.

Figure 10.3A illustrates the *magnitude response* of a low-pass filter. Signals lower than the cutoff frequency (3 MHz in this case) are passed with some small amount of attenuation while signals higher than that frequency are attenuated. The degree of attenuation is dependent on several variables, filter complexity being a major factor.

Of the graphs in this chapter that show a filter's magnitude response, the vertical axes are labeled "Transmission (dB)" with 0 dB at the top of the axis and negative values increasing toward the X-axis. Increasingly negative values of transmission in dB are the same as increasingly positive values of attenuation in dB. For example, –40 dB transmission is the same as 40 dB of attenuation.

Figure 10.3B illustrates the magnitude response of a high-pass filter. Signals above the 3 MHz cutoff frequency are passed with minimum attenuation while signals below that frequency are attenuated. Again, the degree of attenuation is dependent on several variables.

Figure 10.3C illustrates the magnitude response of a band-pass filter. Signals within the band-pass range (between the lower and upper cutoff frequencies) are passed with minimum attenuation while signals outside that range are attenuated. In this example the filter was designed with cutoff frequencies of 2 MHz and 4 MHz, for a passband width of 2 MHz.

Figure 10.3D illustrates the magnitude response of a band-stop filter. Signals within the band-stop range are attenuated while all other signals are passed with minimum attenuation. A notch filter is a type of band-stop filter with a narrow stop-band in which the attenuation is a maximum at a single frequency.

An *ideal filter* — a low-pass filter, for example — would pass all frequencies up to some point with no attenuation at all and totally reject everything beyond that point. This is known as a *brick wall* response because the filter's passband and stopband are *flat*, meaning no attenuation, and the rolloff in the transition region is infinitely steep. The magnitude response of such a filter would be drawn as a flat line representing 0 dB attenuation up to the cutoff frequency that then abruptly changes at the cutoff frequency to a flat line at infinite attenuation throughout the stopband to infinite frequency.

Figure 10.4 shows all of the basic para-

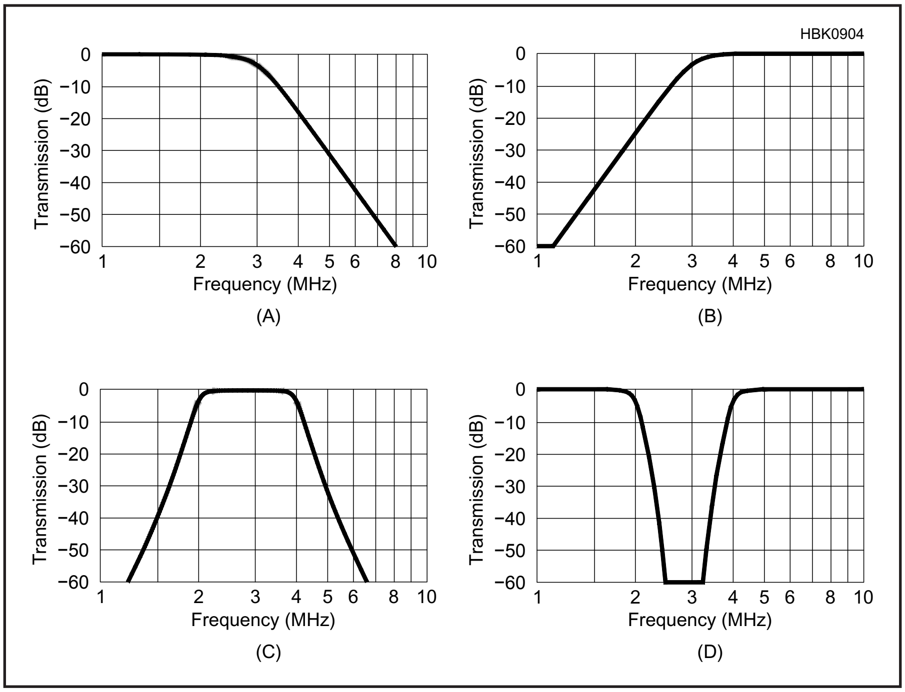

Figure 10.3 — Examples of a low-pass magnitude response (A), high-pass magnitude response (B), band-pass magnitude response (C), and band-stop magnitude response (D).

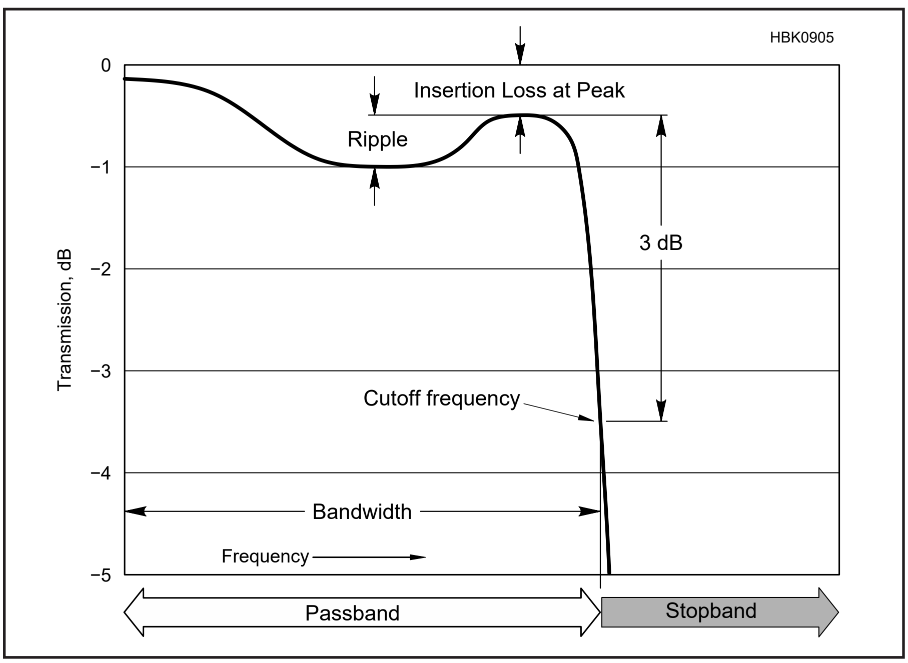

Figure 10.4 — Low-pass filter characteristics showing the passband and stopband, bandwidth, 3 dB cutoff, passband ripple, and insertion loss (IL). This filter has approximately 0.5 dB IL at the frequency of peak response while passband ripple is also 0.5 dB. The vertical axis shows gain (or loss) through the filter, assuming both the input and output are properly terminated. The horizontal axis represents frequency.

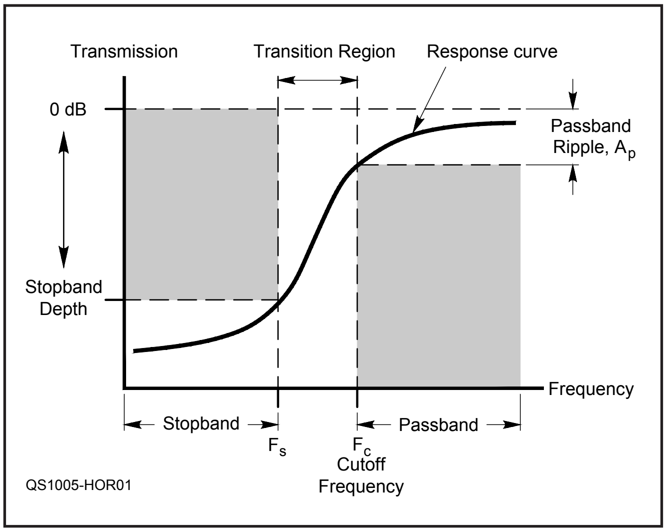

Figure 10.5 — The key specifications for a filter's amplitude response. As long as the filter response curve stays between the light gray boxes, the filter meets the design specification.

meters for a low-pass filter. Variations in amplitude of the response curve in the passband are called *ripple*. (Filters can also have ripple in the stopband for some design families as discussed later.) If the filter is made from passive components, there will be some losses in those components — the amount of loss is called *insertion loss* (IL). The attenuation for signal frequencies in the stopband far from the cutoff frequency is the filter's *ultimate attenuation*.

SPECIFYING A FILTER RESPONSE

Figure 10.5 shows a high-pass filter's *response curve* and some of the terms used to specify the filter's design. On the frequency axis, F_C is the cutoff frequency and F_S defines the *stopband width*. The passband for a high-pass filter consists of all frequencies above F_C. The stopband for a high-pass filter consists of all frequencies below F_S. The *transition region* is the range of frequencies between the stopband and passband. (For a low-pass filter, F_S and F_C are reversed along with the gray boxes they bound.)

The gray boxes leave a space between them through which the filter's response curve must pass. The borders of the gray boxes establish the required performance for the filter. The gray box defining the stopband is bounded on the bottom by the stopband depth. The filter's transmission in the stopband must be equal to or below the stopband depth. The gray box defining the passband is bounded on the top by the passband ripple. The filter's transmission in the passband must be equal to or greater than the passband ripple. Between F_S and F_C is the transition region in which the response curve passes between the stopband and passband. Any response curve that passes through the space between the gray boxes meets the performance requirements for the filter.

As frequency increases through the transition region from the passband into the stopband, the attenuation increases. *Rolloff*, the rate of change of the attenuation for frequencies above cutoff, is expressed in terms of dB of change per octave of frequency. A filter's response with high values of rolloff is called "sharp" or "steep" and "soft" or "shallow" if rolloff is low.

10.2.2 Filter Order

The steepness of the descent from the passband to the region of attenuation — the stopband — is dependent on the complexity of the filter, called the *order.* In a lumped-element filter made from inductors and capacitors, the order is determined by the number of separate energy-storing elements (either L or C) in the filter. For example, an RC filter with one resistor and one capacitor has an order of 1. A notch filter made of a single series-LC circuit has an order of 2. The definition of order for active and digital signal processing filters can be more complex, but the general understanding remains valid that the order of a filter determines how rapidly its response can change with frequency.

For example, **Figures 10.6** and **10.7** shows the magnitude response of one type of low-pass filter (Butterworth family) with orders varying from very simple ("N=2") to the more complex ("N=20"). For each order, frequency has been *normalized* to the ratio of frequency, f, to cutoff frequency, f_C. (i.e. Normalized cutoff frequency is always 1.0.) As the order increases, the rolloff also increases. Figure 10.6 has an expanded vertical scale to show the filter's behavior in the passband more clearly.

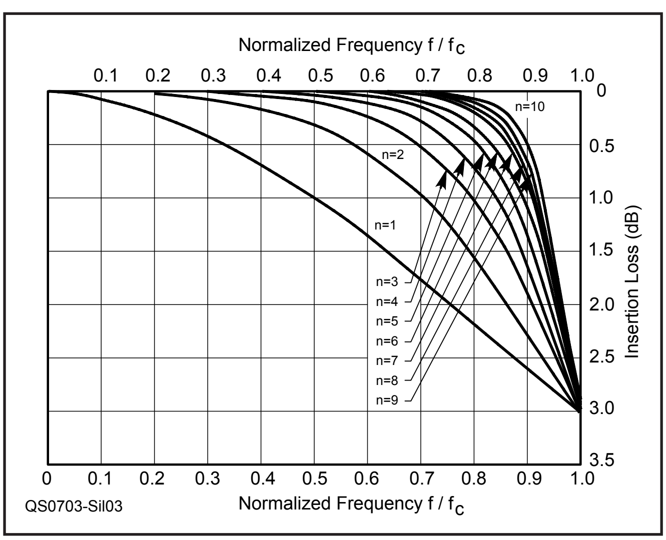

Figure 10.6 — Response of Butterworth filters in the passband below f_C. Frequency is shown as the normalized frequency, f/f_C.

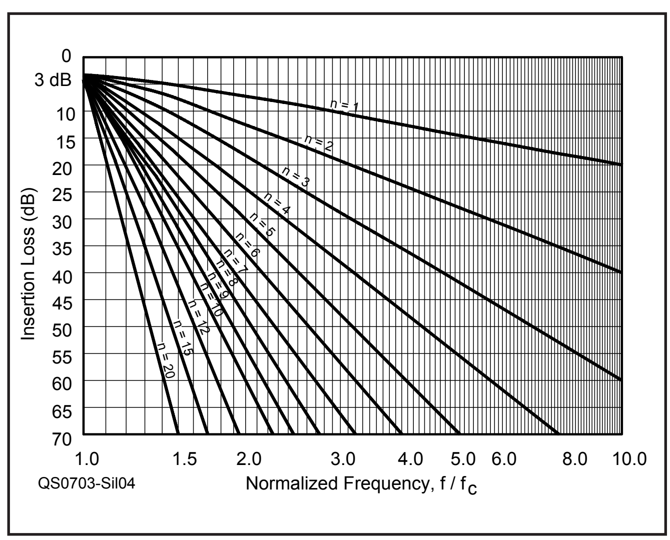

Figure 10.7 — Response of Butterworth filters in the transition and stop-bands above f_C. Frequency is shown as the normalized frequency, f/f_C.

10.2.3 Filter Families

There is no single "best" way to design a filter. Instead we have to decide on some traits and then choose the most appropriate *family* for our design. Different families have different traits, and filter families are commonly named after the mathematician or engineer responsible for defining their behavior mathematically. Each filter family is represented by a specific type of equation that describes the filter's behavior.

Two primary traits of the most importance to amateurs are used to describe the behavior of filter families: ripple (variations in the magnitude response within the passband and stopband) and rolloff. Different families of filters have different degrees of ripple and rolloff (and other characteristics, such as phase response).

With a flat magnitude response (and so no ripples in the passband) we have what is known as a *Butterworth* family design. This family also goes by the name of *Maximally Flat Magnitude* or *Maximally Flat Gain*. An example of the magnitude response of a filter from the Butterworth family is shown in **Figure 10.8**. (Most filter discussions are based on low-pass filters because the same concepts are easily extended to other types of filters and much of the mathematics behind filters is equivalent for the various types of frequency responses.)

By allowing magnitude response ripples in the passband, we can get a somewhat steeper rolloff from the passband into the stopband, particularly just beyond the cutoff frequency. A family that does this is the *Chebyshev* family. **Figure 10.9** illustrates how allowing magnitude ripple in the passband provides a sharper filter rolloff. This plot compares the 1-dB ripple Chebyshev with the no-ripple Butterworth filter down to 12 dB of attenuation.

Figure 10.9 also illustrates the usual definition of bandwidth for a low-pass Butterworth filter — the filter's 3-dB or cutoff frequency. For a low-pass Chebyshev filter, *ripple bandwidth* is used — the frequency range over which the filter's passband ripple is no greater than the specified limit. For example, the ripple bandwidth of a low-pass Chebyshev filter designed to have 1 dB of passband ripple is the highest frequency at which attenuation is 1 dB or less. The 3-dB bandwidth of the frequency of the filter will be somewhat greater.

A Butterworth filter is defined by specifying the order and bandwidth. The Chebyshev filter is defined by specifying the order, the ripple bandwidth, and the amount of passband ripple. In Figure 10.9, the Chebyshev filter has 1 dB of ripple; its ripple bandwidth is 1000 Hz. The Butterworth filter has a 3-dB bandwidth also of 1000 Hz. Some filter textbooks use the 3-dB point to define Chebyshev filters; most use the ripple bandwidth as illustrated here. The schematics (if you ignore parts values) of those two families are identical.

Even small amounts of ripple can be beneficial in terms of increasing a filter's rolloff. **Figure 10.10** compares a Butterworth filter (with the narrow line plot, no ripple in the passband) with a Chebyshev filter (wide line plot, 0.2 dB of ripple in the passband) down to 60 dB of attenuation. (For this comparison the cutoff frequencies at 3 dB of attenuation are the same for each filter.) Even that small amount of ripple in the Chebyshev filter passband allows a noticeably steeper rolloff between the passband and the stopband. As the passband ripple specification is increased, the steepness of the transition from passband to stopband increases, compared to the Butterworth family although the rolloff of the two families eventually becomes equal.

For Chebyshev filters, when the value of the passband ripple is changed, the magnitude response in the stopband region also changes. **Figure 10.11** compares the stopband response of Chebyshev filters with passband ripple ranging from 0.01 to 1 dB.

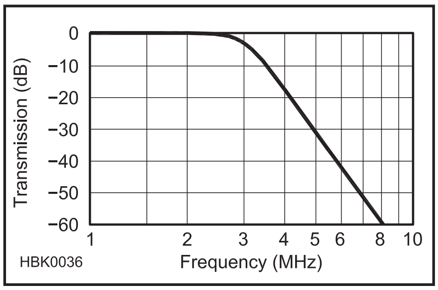

Figure 10.8 — Filters from the Butterworth family exhibit flat magnitude response in the passband.

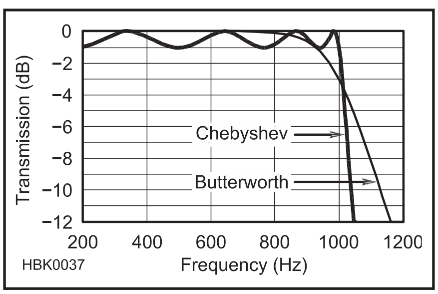

Figure 10.9 — This plot compares the response of a Chebyshev filter with a 1-dB ripple bandwidth of 1000 Hz and a Butterworth filter 3-dB bandwidth of 1000 Hz.

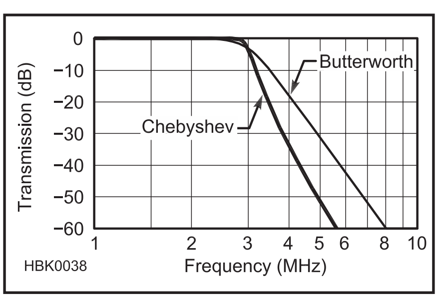

Figure 10.10 — A Chebyshev filter (0.2 dB passband ripple) allows a sharper cutoff than a Butterworth design with no passband ripple.

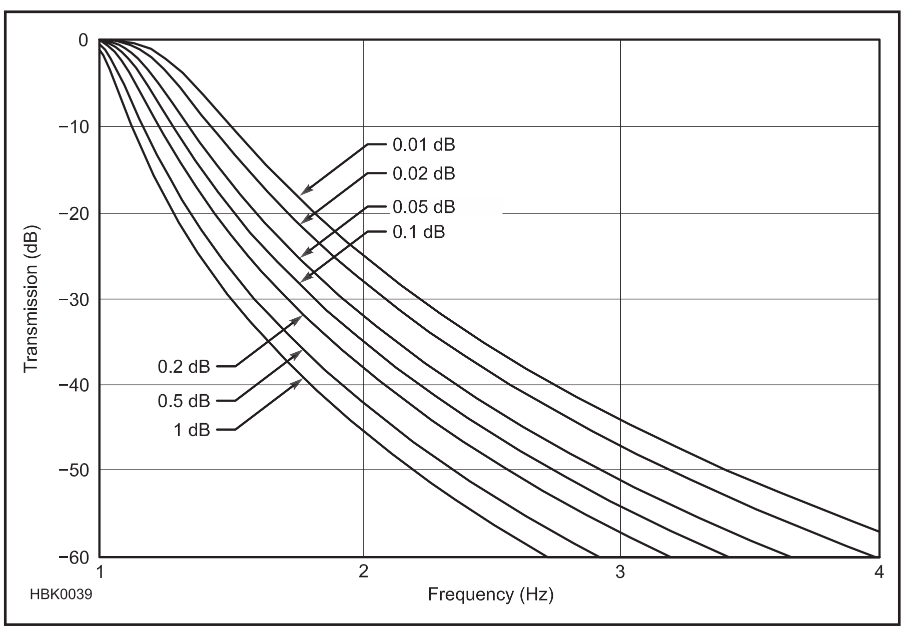

Figure 10.11 — These stop band response plots illustrate the Chebyshev family with various values of passband ripple. These plots are for a seventh-order low-pass design with ripple values from 0.01 to 1 dB. Ignoring the effects in the passband of high ripple values, increasing the ripple will allow somewhat steeper rolloff into the stop band area, and better ultimate attenuation in the stop band.

It is possible to obtain even steeper rolloff into the stopband by adding "traps" whose frequencies are carefully calculated. The resonant frequencies of those traps are in the stopband region and are set to yield best performance. When this is done, we have a *Cauer* family design (also called the *elliptic-function* design). The rolloff of the Cauer filter from the passband into the stopband is the steepest of all analog filter types provided that the behavior in the passband is uniform (either no ripple or a uniform amount of ripple). **Figure 10.12** shows the response of the Chebyshev and Cauer designs for comparison.

The Cauer filter is defined by specifying the order, the ripple bandwidth, and the passband ripple, just as for the Chebyshev. Again, an alternative bandwidth definition is to use the 3-dB frequency instead of the ripple bandwidth. The Cauer family requires one more specification: the *stopband frequency* and/or the *stopband depth*. The stopband frequency is the lowest frequency of a null or notch in the stopband. The stopband depth is the minimum amount of attenuation allowed in the stopband. In Figure 10.12, the stopband frequency is about 3.5 MHz and the stopband depth is 50 dB. For this comparison both designs have the same passband ripple value of 0.2 dB.

A downside to the Cauer filter is that the ultimate attenuation is some chosen value rather than ever increasing, as is the case with the other families. In the far stopband region (where frequency is much greater than the cutoff frequency) the rolloff of Cauer filters ultimately reaches 6 or 12 dB per octave, depending on the order. Odd-ordered Cauer filters have an ultimate rolloff rate of about 6 dB per octave while the even-ordered versions have an ultimate rolloff rate of about 12 dB per octave.

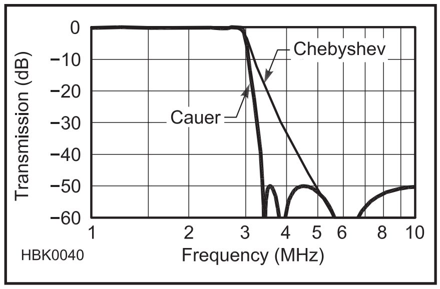

Figure 10.12 — The Cauer family has an even steeper rolloff from the passband into the stop band than the Chebyshev family. Note that the ultimate attenuation in the stop band is a design parameter rather than ever-increasing as is the case with other families.

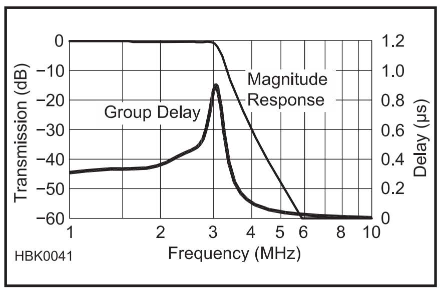

Figure 10.13 — Magnitude response and group delay of a Chebyshev low-pass filter.

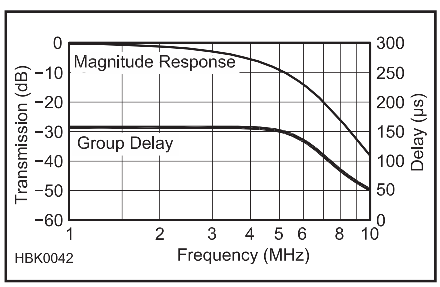

Figure 10.14 — Magnitude response and group delay of a Bessel low-pass filter.

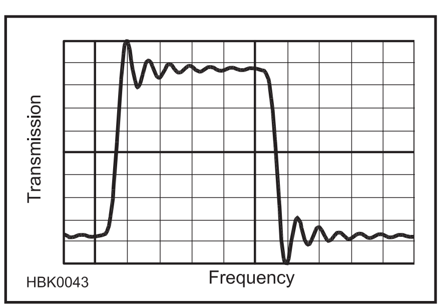

Figure 10.15 — Transient response of a Chebyshev low-pass filter.

10.2.4 Group Delay

Another trait that can influence the choice of which filter family to use is the way the transit time of signals through the filter varies with frequency. This is known as *group delay*. ("Group" refers to a group of waves of similar frequency and phase moving through a media, in this case, the filter.) The wide line (lower plot) in **Figure 10.13** illustrates the group delay characteristics of a Chebyshev low-pass filter, while the upper plot (narrow line) shows the magnitude response. As shown in Figure 10.13, the group delay of components near the cutoff frequency becomes quite large when compared to that of components at lower and higher frequencies. This is a result of the phase shift of the filter's transmission being nonlinear with frequency; it is usually greater near the cutoff frequency. By delaying signals at different frequencies different amounts, the signal components are "smeared" in time. This causes distortion of the signal and can seriously disrupt high-speed data signals.

If a uniform group delay for signals throughout the passband is needed, then the *Bessel* filter family should be selected. The Bessel filter can be used as a delay-line or time-delay element although the gentle rolloff of the magnitude response may need to be taken into account. The magnitude and delay characteristics for the Bessel family are shown in **Figure 10.14**. Comparing the Bessel filter's response to that of the Butterworth in Figure 10.8 shows the difference in roll off.

A downside of Bessel family filters is that the rolloff characteristic is quite poor; they are not very good as a magnitude-response-shaping filter. The Bessel family is characterized largely by its constant group delay in the passband (for a low- or high-pass filter) shown as the bottom plot in Figure 10.14. The *constant-delay* characteristic of the Bessel extends into the stopband. The Bessel filter bandwidth is commonly defined by its 3-dB point (as with the Butterworth).

Because the Bessel filter is used when phase response is important, it is often characterized by the frequency at which a specific amount of phase shift occurs, usually one radian. (One radian is equal to $360 / 2\pi = 57.3°$.) Since the delay causes the output signal to lag behind the input signal, the filter is specified by its *one-radian lag frequency*.

10.2.5 Transient Response

Some applications require that a signal with sharp rising and falling edges (such as a digital data waveform) applied to the input of a low-pass design have a minimum *overshoot* or *ringing* as seen at the filter's output. Overshoot (and undershoot) occurs when a signal exceeds (falls below) the final amplitude temporarily before settling at its final value. Ringing is a repeated sequence of overshoot and undershoot.

If a sharp-cutoff analog filter is used in such an application, overshoot or ringing will occur. The appearance of a signal such as a square wave with sharp rising and falling edges as it exits from the filter may be as shown in **Figure 10.15**. The scales for both the X- and the Y-axes would depend on the frequency and magnitude of the waveform. (As described later in this chapter, digital filters can be designed not to exhibit this ringing, although they are much more complex than simple analog filters and have other tradeoffs.)

The sharp edge of the square wave is a type of *transient*, an abrupt change in a signal. When a signal at a constant level and changes to another level very rapidly where it remains constant, that type of transient is a *step*. If the signal abruptly changes levels and immediately changes back again, that type of transient is a *pulse*. If the pulse is infinitely narrow, it is called an *impulse* and has interesting mathematical characteristics as discussed below. The output of a circuit a transient occurs at its input is the circuit's *transient response*.

The square wave used in this discussion as a test waveform is composed of a fundamental and an endless series of odd harmonics. If harmonics only up to a certain order are used to create the square wave — that is, if the square wave is passed through a sharp-cutoff low-pass filter that attenuates higher frequencies — then that waveform will have the overshoot or ringing as shown as it exits from the filter. This is the filter's *step response*.

Transients can be repetitive, such as the edges of the square wave, but even a non-periodic waveform can be decomposed into sine waves, although in this case they are not harmonically-related. For example a single pulse of width τ seconds has a frequency spectrum proportional to sinc(fτ) = sin(πfτ)/(πfτ). You can think of this as an infinite number of sine waves spaced infinitely closely together with amplitudes that trace out that spectral shape. It is interesting to note that if τ is decreased, the value of f must increase by the same factor for any given value of sin(πfτ)/(πfτ). In other words, the narrower the pulse the wider the spectrum. Of course that applies to sine waves and other periodic waveforms as well — the smaller the wavelength the higher the frequency. In general, anything that makes the signal "skinnier" in the time domain makes it "fatter" in the frequency domain and vice versa.

As the pulse becomes narrower and narrower, the frequency spectrum spreads out more and more. In the limit, if the pulse is made infinitely narrow (an impulse), the spectrum becomes flat from zero hertz to infinity. The impulse is a very useful concept because of its flat frequency spectrum. The filter's *impulse response* has a frequency spectrum equal to the frequency response of the filter.

10.2.6 Filter Family Selection

Selecting a filter family is one of the first steps in filter design. To make that choice easier, the following list of filter family attributes is provided:

• *Butterworth* — No ripple in passband, smooth transition region, shallow rolloff for a given filter order, high ultimate attenuation, smooth group delay change across transition region. The smoothness of the response is particularly apparent near dc for the low-pass response, at the center frequency for a band-pass response, and at infinity for the high-pass response. The resulting magnitude response will also have a relatively gentle transition from the passband into the stop band.

• *Chebyshev* — Some passband ripple, abrupt transition region, steep rolloff for a given filter order, peak in group delay near cutoff frequency, high ultimate attenuation. The Chebyshev family is used when a sharper cutoff is desired for a given number of components and where at least a small amount of ripple is allowable in the passband.

• *Cauer* (or *Elliptical-function*) — Some passband ripple, abrupt transition region, steepest roll off, ripple in stopband due to traps, ultimate attenuation smaller than Butterworth and Chebyshev, group delay peaks near cutoff frequency and in stopband. When steepness of rolloff from passband into stop band is the item of greatest importance, then the Cauer filter family is used. Cauer filters involve a more complicated set of choices. In addition to selecting a passband ripple, the designer must also assign a stop band depth (or stop band frequency). Some of the items interact; they can't all be selected arbitrarily.

• *Bessel* — No ripple in passband, smooth transition region, constant group delay in passband, shallow rolloff for a given filter order, smooth changes in group delay, high ultimate attenuation

All characterizations such as "steep" and "abrupt" are relative with respect to filter designs from other families with similar orders. Other factors, such as number of components, sensitivity to component value and so on may need to be considered when selecting a filter family for a specific application.

All of the traits mentioned so far in this chapter apply to a filter regardless of how it is implemented, whether it is fabricated using *passive* lumped-element inductors and capacitors or using op amps with resistors and capacitors (an *active* filter) or in software as a digital filter. The traits are general descriptions of filter behavior and can be applied to any type of filter technology. Each type of technology (passive, active, digital) has strong points and tradeoffs.

These four families are the most common, but other families are used as well: for example, Gaussian, Constant-k, and M-derived are all supported by the *ELSIE* filter design package provided with this book's online information

Digital filters also implement these filter families and others, as well. They also have other classes, such as FIR and IIR, as described in this chapter's section on digital filters. For more information about analog and digital filter design, a list of references and articles is provided at the end of this chapter.

10.3 Passive LC Filters

This part of the chapter deals with passive LC filters fabricated using discrete inductors and capacitors (which gives rise to their name, *lumped-element*). We will begin with a discussion of basic low-pass filters and then generalize to other types of filters.

10.3.1 Low-Pass Filters

A very basic LC filter built using inductors and capacitors is shown in **Figure 10.16**. In the first case (Figure 10.16A), less power is delivered to the load at higher frequencies because the reactance of the inductor in series with the load increases as the test frequency increases. The voltage appearing at the load goes down as the frequency increases. This configuration would pass direct current (dc) and reject higher frequencies, and so it would be a low-pass filter.

With the second case (Figure 10.16B), less power is delivered to the load at higher frequencies because the reactance of the capacitor in parallel with the load decreases as the test frequency increases. Again, the voltage appearing at the load goes down as the frequency increases and so this, too, would be called a low-pass filter. In the real world, combinations of both series and parallel components are used to form a low-pass filter.

A high-pass filter can be made using the opposite configuration — series capacitors and shunt inductors. And a band-pass (or band-stop) filter can be made using pairs of series and parallel tuned circuits. These filters, made from alternating LC elements or LC tuned circuits, are called *ladder filters*.

We spoke of filter order, or complexity, earlier in this chapter. **Figure 10.17** illustrates *capacitor-input* low-pass filters with orders of 3, 4 and 5. Remember that the order corresponds to the number of energy-storing elements. For example, the third-order filter in Figure 10.17A has three energy-storing elements (two capacitors and one inductor), while the fifth-order design in Figure 10.17C has five elements total (three capacitors, two inductors). For comparison, a third-order filter is illustrated in **Figure 10.18**. The filters in Figure 10.17 are *capacitor-input* filters because a capacitor is connected directly across the input source. The filter in Figure 10.18 is an *inductor-input* filter.

As mentioned previously, the Cauer family has traps (series or parallel tuned circuits)

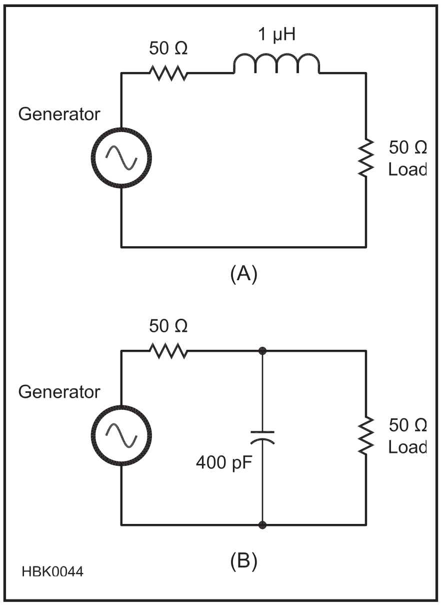

Figure 10.16 — A basic low-pass filter can be formed using a series inductor (A) or a shunt capacitor (B).

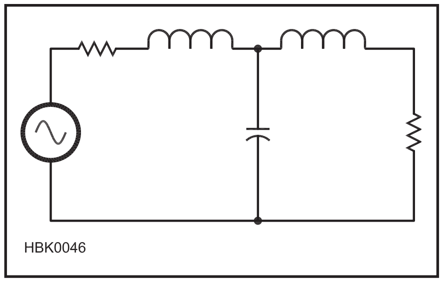

Figure 10.18 — A third-order inductor-input low-pass filter.

carefully added to produce dips or notches (properly called *zeros*) in the stop band. Schematics for the Cauer versions of capacitor-input low-pass filters with orders of 3, 4 and 5 are shown in **Figure 10.19**. The capacitors in parallel with the series inductors create the notches at calculated frequencies to allow the Cauer filter to be implemented.

The capacitor-input and the inductor-input versions of a given low-pass design have identical characteristics for their magnitude, phase and time responses, but they differ in the impedance seen looking into the filter. The capacitor-input filter has low impedance in the stop band while the inductor-input filter has high impedance in the stop band.

10.3.2 Low-Pass to Band-Pass Transformation

A band-pass filter is defined in part by a *bandwidth* and a *center frequency*. (An alternative method is to specify a lower and an upper cutoff frequency.) A low-pass design such as the one shown in Figure 10.17A can be converted to a band-pass filter by resonating each of the elements at the center frequency. **Figure 10.20** shows a third-order low-pass filter with a design bandwidth of 2 MHz for use in a 50 Ω system. If the shunt elements are now resonated with a parallel component, and if the series elements are resonated with a series component, the result is a band-pass filter as shown in **Figure 10.21**.

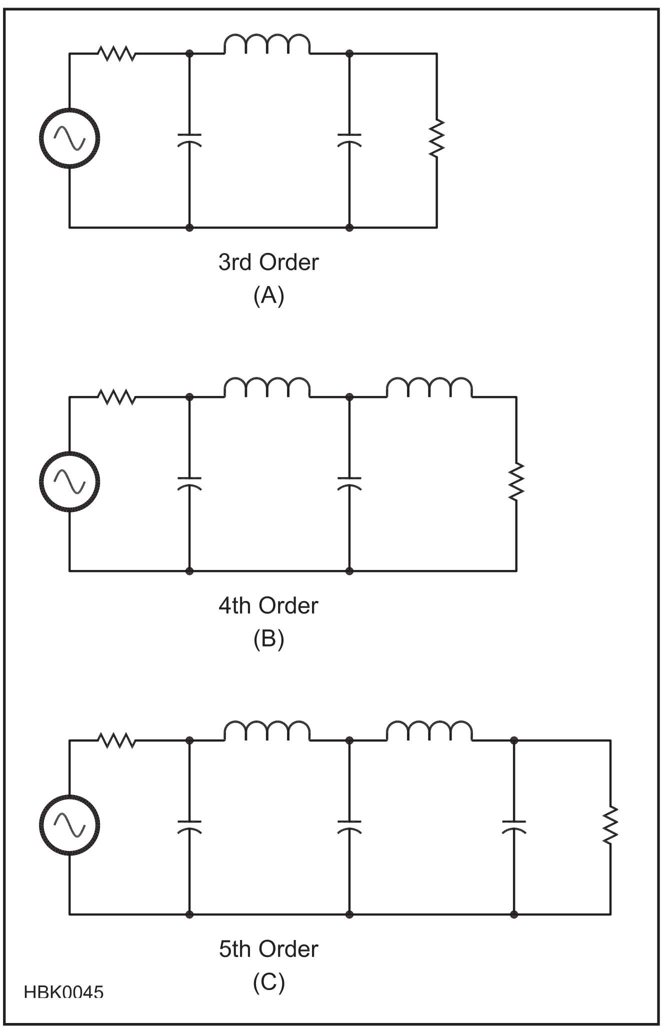

Figure 10.17 — Low-pass, capacitor-input filters for the Butterworth and Chebyshev families with orders 3, 4 and 5.

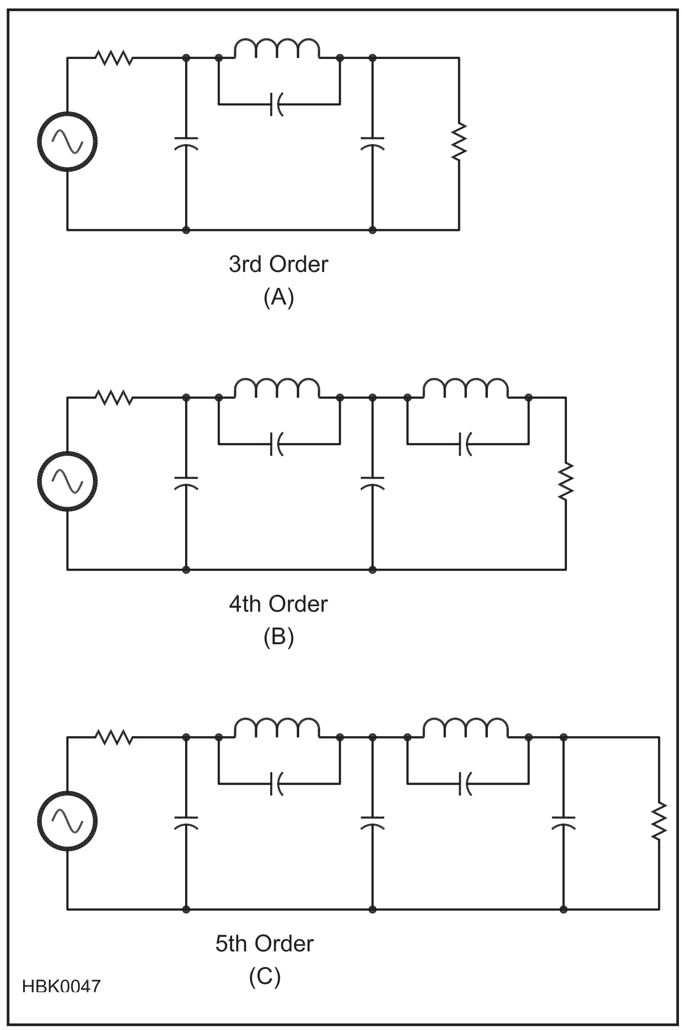

Figure 10.19 — Low-pass topologies for the Cauer family with orders 3, 4 and 5.

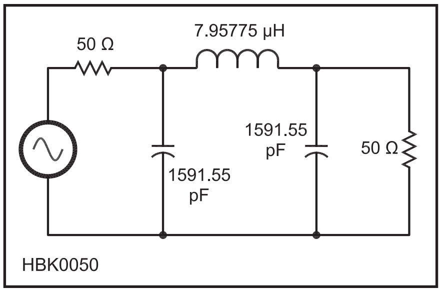

Figure 10.20 — Third-order low-pass filter with a bandwidth of 2 MHz in a 50 Ω system.

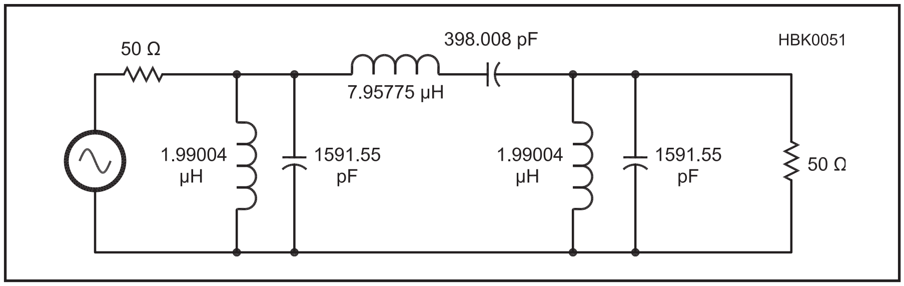

Figure 10.21 — The low-pass filter of Figure 10.20 can be transformed to a band-pass filter by resonating the shunt capacitors with a parallel inductor and resonating the series inductor with a series capacitor. Bandwidth is 2 MHz and the center frequency is 2.828 MHz.

The series inductor value and the shunt capacitor values are the same as those for the original low-pass design.

10.3.3 High-Pass Filters

A high-pass filter passes signals above its cutoff frequency and attenuates those below. Simple high-pass equivalents of the filters in Figure 10.16 are shown in **Figure 10.22.**

The reactance of the series capacitor in Figure 10.22A increases as the test frequency is lowered and so at lower frequencies there will be less power delivered to the load. Similarly, the reactance of the shunt inductor in Figure 10.22B decreases at lower frequency, with the same effect. Similarly to the low-pass filter designs presented in Figure 10.17, a high-pass filter in a real world design would typically use both series and shunt components but with the positions of inductors and capacitors exchanged.

An example of a high-pass filter application would be a broadcast-reject filter designed to pass amateur-band signals in the range of 3.5 MHz and above while rejecting broadcast signals at 1.7 MHz and below. A high-pass filter with a design cutoff of 2 MHz is illustrated in **Figure 10.23**. It can be implemented as a capacitor-input (Figure 10.23A) or inductor-input (Figure 10.23B) design. In each case, signals above the cutoff are passed with minimum attenuation while signals below the cutoff are attenuated, in a manner similar to the action of a low-pass filter.

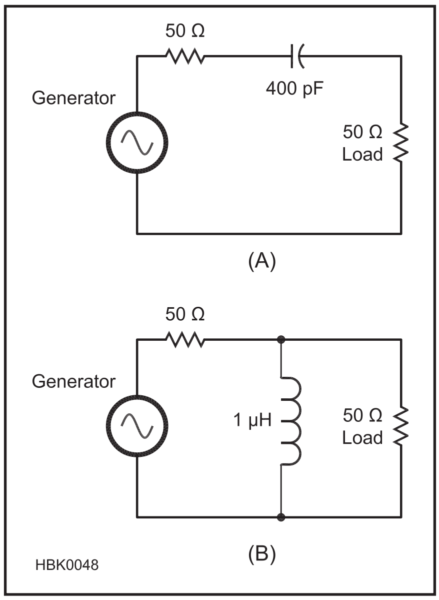

Figure 10.22 — A basic high-pass filter can be formed using a series capacitor (A) or a shunt inductor (B).

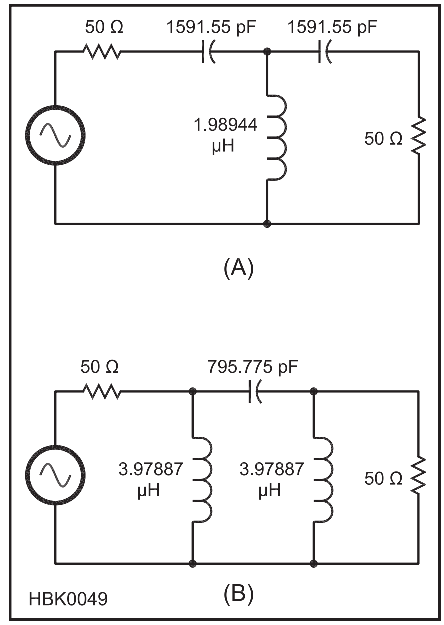

Figure 10.23 — Capacitor-input (A) and inductor-input (B) high-pass filters. Both designs have a 2 MHz cutoff.

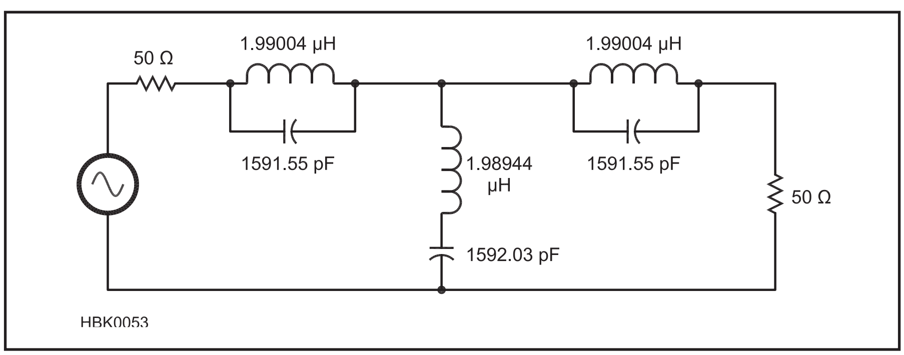

Figure 10.24 — The high-pass filter of Figure 10.23A can be converted to a band-stop filter with a bandwidth of 2 MHz, centered at 2.828 MHz.

10.3.4 High-Pass to Band-Stop Transformation

Just as a low-pass filter can be transformed to a band-pass type, a high-pass filter can be transformed into a band-stop (also called a *band-reject*) filter. The procedures for doing this are similar in nature to those of the transformation from low-pass to band-pass. As with the band-pass filter example, to transform a high-pass to a band-stop we need to specify a center frequency. The bandwidth of a band-stop filter is measured between the frequencies at which the magnitude response drops 3 dB in the transition region into the stop band.

Figure 10.24 shows how to convert the 2 MHz capacitor-input high-pass filter of Figure 10.20A to a band-stop filter centered at 2.828 MHz. The original high-pass components are resonated at the chosen center frequency to form a band-stop filter. Either a capacitor-input high-pass or an inductor-input high-pass may be transformed in this manner. In this case the series capacitor values and the shunt inductor values for the band-stop are the same as those for the high-pass. The series elements are resonated with an element in parallel with them. Similarly, the shunt elements are also resonated with an element in series with them. In each case the pair resonate at the center frequency of the band-stop.

10.3.5 Effect of Component Q

When components with less-than-ideal characteristics are used to fabricate a filter, the performance will also be less than ideal. One such item to be concerned about is component "Q." As described in the **Radio Fundamentals** chapter, Q is a measure of the loss in an inductor or capacitor, as determined by its resistive component. Q is the ratio of component's reactance to the loss resistance and is specified at a given test frequency. The loss resistance referred to here includes not only the value as measured by an ohmmeter but includes all sources of loss, such as skin effect, dielectric heating and so on.

The Q values for capacitors are usually greater than 500 and may reach a few thousand. Q values for inductors seldom reach 500 and may be as low as 20 or even worse for miniaturized parts. A good toroidal inductor can have a Q value in the vicinity of 250 to 400.

Q values can affect both the *insertion loss* of signals passing through the filter and the steepness of the filter's rolloff. Band-pass filters (and especially narrowband band-pass filters) are more vulnerable to this problem than low-pass and high-pass filters. **Figure 10.25** shows the effect of finite values of inductor Q values on the response of a low-pass filter. The Q values for each plot are as shown.

Inadequate component Q values introduce loss and more importantly they compromise the filter's response at cutoff, especially problematic in the case of narrowband band-pass filters. **Figure 10.26** illustrates the effect of finite inductor Q values on a narrowband band-pass filter. In the case of a band-pass filter, the Q values required to support a given response shape are much higher than those required for the low-pass or high-pass filter (by the ratio of center to width). Capacitor Q values are generally much high than inductor Q values and so contribute far less to this effect.

In general, component-value adjustment will not be able to fully compensate for inadequate component Q values. However, if the filter is deliberately mismatched (by changing the input and/or output terminations) then a limited amount of *response-shape* correction can sometimes be achieved by network component value optimization ("tweaking"). The loss caused by Q problems (at dc in the case of a low-pass or at the center frequency in the case of a band-pass) may increase if such correction is attempted.

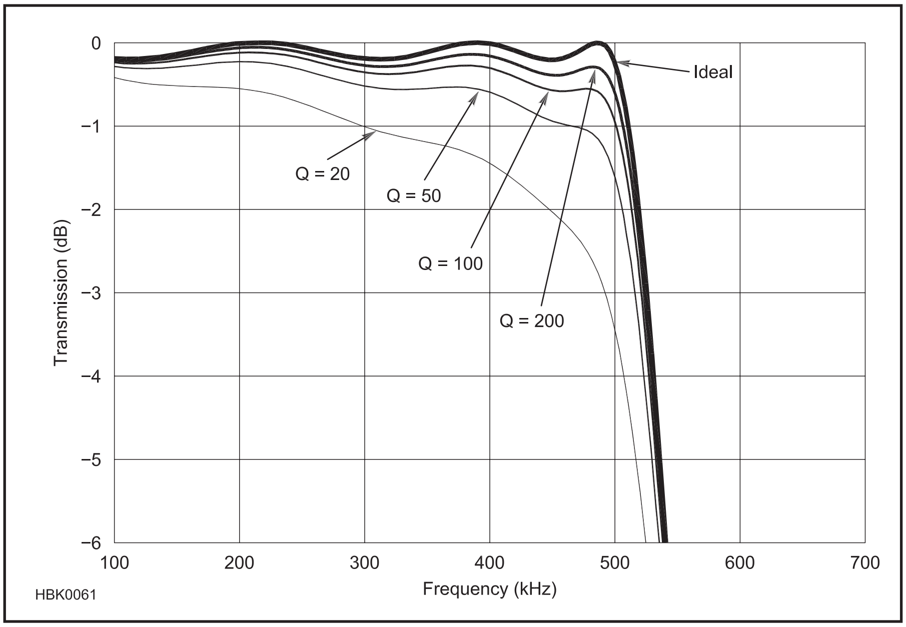

Figure 10.25 — Effect of inductor Q values on a low-pass filter.

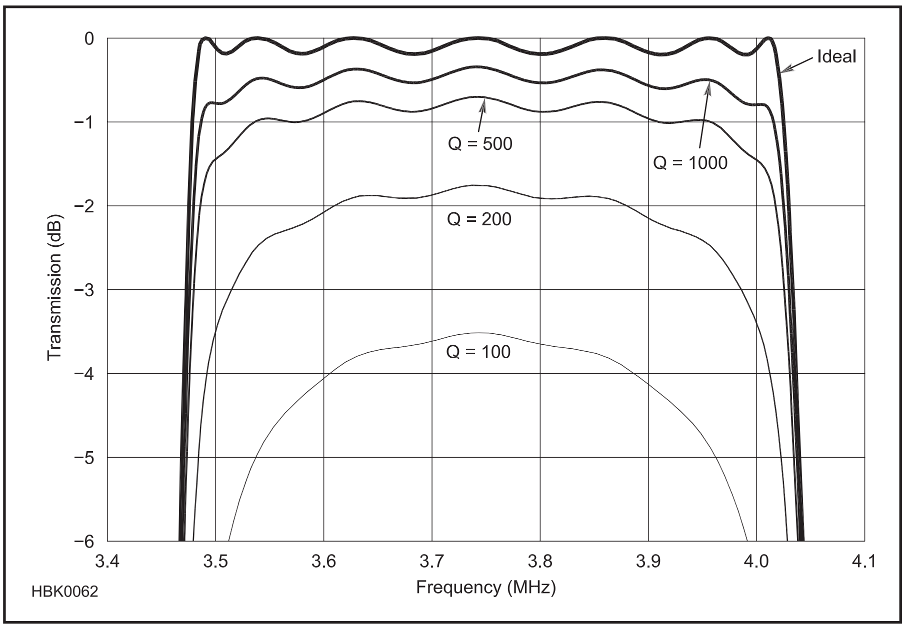

Figure 10.26 — Effect of inductor Q values on a narrowband band-pass filter.

10.3.6 Side Effects of Passband Ripple

Especially in RF applications it is desirable to design a filter such that the impedance seen looking into the input side remains fairly constant over the passband. Increasing the value of passband ripple increases the rate of descent from the passband into the stop band, giving a sharper cutoff. But it also degrades the uniformity of the impedance across the passband as seen looking into the input of the filter. This may be shown in terms of VSWR or return loss; those are simply different ways of stating the same effect. (*Return loss* is explained in the **RF Techniques** chapter.)

Designs using a low value of passband ripple are preferred for RF work. Audio-frequency applications are generally not as critical, and so higher ripple values (up to about 0.2 dB) may be used in audio work.

Table 10.2 shows the maximum value of VSWR and minimum value of return loss for various values of passband ripple. Note that lower values for passband ripple yield better

Table 10.2
Passband Ripple, VSWR and Return Loss

Passband Ripple (dB)	*VSWR*	*Return Loss (dB)*
0.0005	1.022	39.38
0.001	1.031	36.37
0.002	1.044	33.36
0.005	1.07	29.39
0.01	1.101	26.38
0.02	1.145	23.37
0.05	1.24	19.41
0.1	1.355	16.42
0.2	1.539	13.46
0.5	1.984	9.636
1	2.66	6.868

Note: As the passband ripple specification is changed so do the other items. Conversely, to limit VSWR to a maximum value or return loss to at least some minimum value, use this table to find the passband ripple that should be used to design the filter.

values for VSWR and return loss. Specifying a filter to have a passband ripple value of 0.01 dB will result in that filter's VSWR figure to be about 1.1:1. Or restated, the return loss will be about 26 dB. These values will be a function of frequency and at some test frequencies may be much better.

10.3.7 Use of Filters at VHF and UHF

Even when filters are designed and built properly, they may be rendered totally ineffective if not installed properly. Leakage around a filter can be quite high at VHF and UHF, where wavelengths are short. Proper attention to shielding and good grounding is mandatory for minimum leakage. Poor coaxial cable shield connection into and out of the filter is one of the greatest offenders with regard to filter leakage. Proper dc-lead bypassing throughout the receiving system is good practice, especially at VHF and above. Ferrite beads placed over the dc leads may help to reduce leakage. Proper filter termination is required to minimize loss.

Most VHF RF amplifiers optimized for noise figure do not have a 50 Ω input impedance. As a result, any filter attached to the input of an RF amplifier optimized for noise figure will not be properly terminated and filter loss may rise substantially. As this loss is directly added to the RF amplifier noise figure, carefully choose and place filters in the receiver.

10.3.8 Design Software for LC Filters

Previous editions of this book included an extensive set of design tables and formulas for manual calculation of filter component values. While this method was certainly instructive, it was tedious and error-prone, particularly for higher-order designs. The sections that described this method in previous editions have been extracted as a PDF document that is provided with the online information for this book.

Filter design is almost universally performed with software today and the *ARRL Handbook* is fortunate to include such a package. Jim Tonne, W4ENE, has made available a version of *ELSIE*, a filter design program, to amateurs at no charge. The latest version of *ELSIE* and related programs are available for downloading with the online content. The list of software is presented at the beginning of this chapter and the program capabilities are listed below.

Users unfamiliar with *ELSIE* will benefit from following the "walkthrough" accessible from the program's ABOUT menu tab. Numerous design examples and tutorials are available as online videos hosted on **YouTube.com** — search for "elsie filter design tutorial" to locate a good starting set. A two-part step-by-step tutorial for using *ELSIE* from the Hands-On Radio series of columns is available in the online material for this book as well.

- *ELSIE* — design and analysis of lumped-element LC filters. In addition to providing parts values for filters with various topologies from various families, tools are included to assist with practical construction.
- *SVC Filter Designer* — design of lumped-element high-pass and low-pass filters. The software shows ideal values and also the nearest 5% values for capacitors and inductors. It also analyzes those filters and shows the deviation of key responses from ideal when those 5% values are used.
- *QuadNet* — design and analysis of active quadrature ("90-degree") networks for use in SSB transmitters and receivers. It handles networks with orders from 2 to 10, odd and even, with tuning modes and analysis.
- *Helical* — design and analysis of helical-resonator band-pass filters usually used in the VHF and UHF frequency ranges.
- *Diplexer* — design and analysis of diplexer filters.

10.4 Active Audio Filters

Below RF, in what is broadly referred to as the "audio" range between a few Hz and a few hundred kHz, designers have several choices of filter technology.

- Passive LC
- DSP Digital Filters
- Switched-Capacitor Audio Filter (SCAF)
- Active RC

LC audio filters are not used much in current designs except in high-power audio applications, such as speaker crossover networks, and are not covered here. LC filters were once popular as external audio filters for CW reception. These designs tended to be large and bulky, often using large surplus 44 or 88 mH core inductors. LC filters used in very low-level receiver applications can be very compact, but these tend to suffer from relatively high insertion loss, which reduces the receiver sensitivity (if no preamp is used ahead of the filter) or its large-signal dynamic range (if a loss-compensating preamp is used ahead of the filter). In addition, LC filters used in low-level receiver applications tend to pick up 60/120 Hz hum from stray magnetic fields from ac power transformers. Even "self shielded" inductors can produce noticeable ac hum pick up when followed by 90 to 120 dB of audio gain!

Digital filtering (see the following section of this chapter) can yield filters that are superior to analog filters. High signal-level DSP-based external audio filters such as the popu-lar MFJ Enterprises MFJ-784B (**www.mfjenterprises.com**) have been available for some time. In addition, if the filter is coupled to a high-performance audio A/D converter, digital filtering can provide excellent filtering for even very low-level audio signals.

The main drawback to standalone digital filters is that of expense. A low-end DSP external audio filter costs at least $100, while a DSP with a high-end A/D converter (such as an audio sound card front-end to a PC) can run in the $150 to $400+ price range, which does not include the cost of the host PC.

10.4.1 SCAF Filters

Simple SCAF filter designs can produce extremely effective low-power audio filters. The implementation of these filters in practical IC form usually involves small-value capacitors and high-value resistors. The use of high-value resistors tends to generate enough noise to make these filters unsuitable for very low-level audio processing such as the front end of a receiver audio chain. Thus, SCAF filters tend to be limited to filtering at the output end of the audio chain — headphone or speaker level audio applications, an area in which they very much excel. The cutoff frequency of a SCAF filter is set by an external clock signal. Thus, this class of filter naturally lends itself to use in variable frequency filters. For simplicity, excellent frequency selectivity and relatively low cost, SCAF filters are highly recommended when additional audio filtering is desired on the output of an existing receiver.

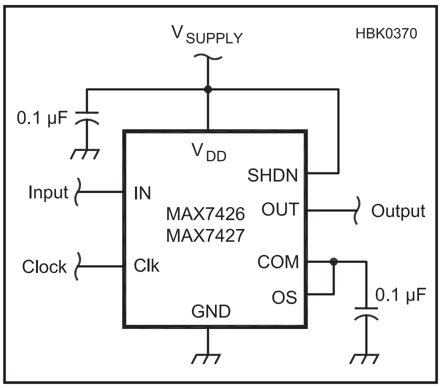

Figure 10.27 — Simple SCAF low pass filter, taken from the MAX7426 data sheet.

An example of a very simple, but highly effective SCAF low-pass filter IC is the Maxim MAX7426 (5 V supply) or MAX7427 (3 V supply) in **Figure 10.27.** The filter's cutoff frequency can be set by placing an appropriately sized capacitor across the clock oscillator inputs because the part can generate its own internal clock signal. For example, connecting a 180 pF capacitor across the MAX7426's clock inputs will produce a 1 kHz low-pass filter. The MAX7427 (3-V version) was used in the NC2030 QRP transceiver along with a MVAM108 varactor diode to create a low-pass filter that was tunable from 300 Hz to 1 kHz.

A portion of that schematic is shown in **Figure 10.28**. The SCAF low-pass filter is followed by a 3-V unity-gain headphone amplifier (U14C and U14D) as the filter IC itself cannot directly drive headphones. The low-pass cutoff frequency is tuned by using R83 to vary the voltage applied across the MVAM108 varactor diode D10, thus changing the capacitance across the clock input (pin 8) of the chip. C116 was used to isolate the dc voltage across the varactor diode from the bias voltage on the clock input line. The frequency response of such a filter is shown in **Figure 10.29**.

As can be seen, this is an extremely sharp filter, producing almost 40 dB of attenuation very close to the cut off frequency — not too bad for an inexpensive part. The slightly more expensive MAX7403 gives an even steeper 80 dB cutoff. Both parts are specified for an 80 dB signal-to-noise ratio based on an assumed 4-V_{P-P} signal. Sensitive modern in-ear type headphones require only 20 mV_{P-P} to produce a fairly loud signal. A 20 mV signal is 46 dB below 4 V_{P-P}, so at such headphone levels, the noise floor is actually only 34 dB below the 4 V_{P-P} signal — fairly quiet, but is a lot less noise margin than one would tend to think given the 80 dB specification.

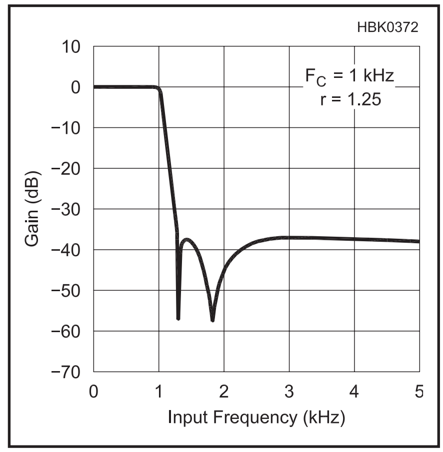

Figure 10.29 — Frequency response of a low-pass filter using the MAX7426 set to a 1 kHz cutoff frequency.

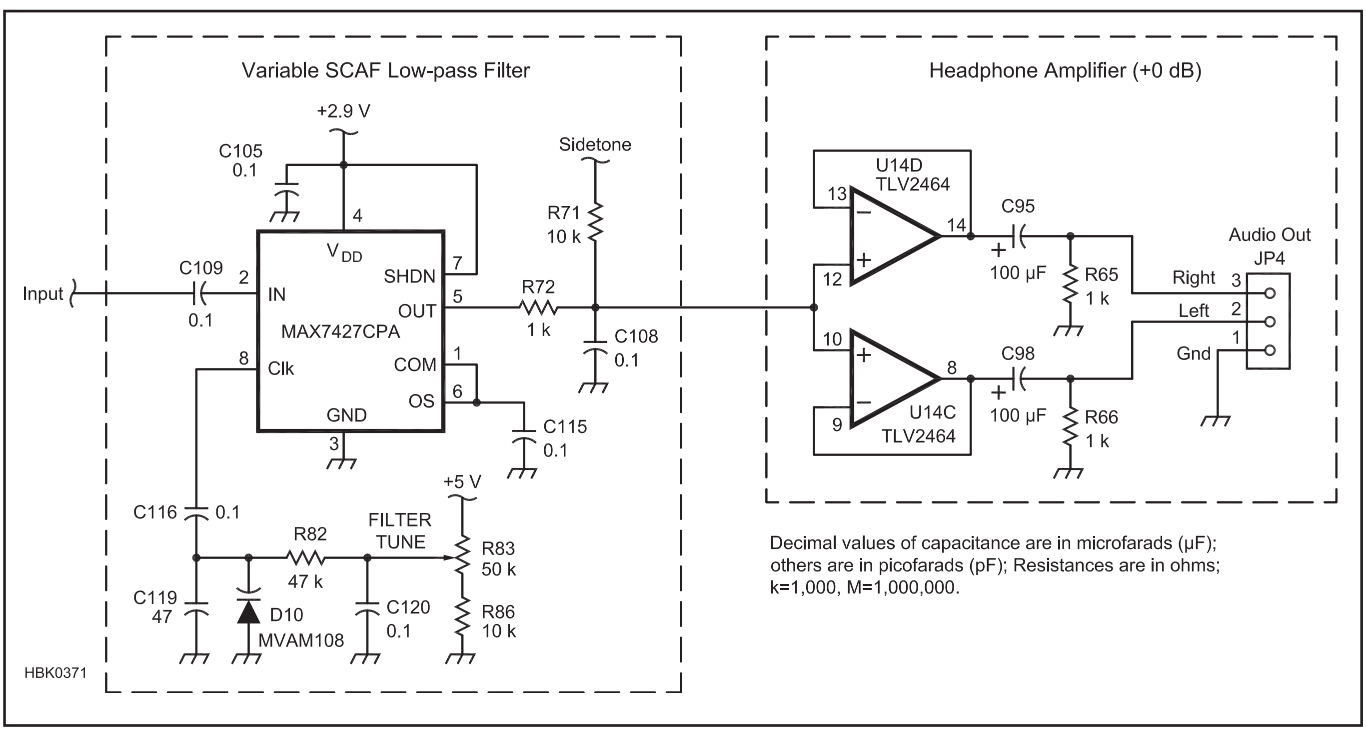

Figure 10.28 — SCAF with variable 300 to 1000 Hz low-pass cutoff and unity-gain headphone amplifier.

Again, with a noise floor only 34 dB below typical headphone levels, SCAF filters such as these are fine at the end of the receiver chain, but are not useful as an audio filter at very low signal levels early on in a receiver.

Figure 10.30 is an example of a slightly more complex SCAF band-pass filter. This filter features both a variable center frequency (450 Hz to 1000 Hz) and a variable bandwidth (90 Hz to 1500 Hz). This very popular filter, the NESCAF, was designed by the New England QRP Club (**www.newenglandqrp.org**) and is currently offered in kit form in batches. Contact the club if you wish to purchase the kit. The design is included here and is well within the range of a home-builder with intermediate construction skills. There are a number of online websites that discuss the filter and how to build it.

The NESCAF filter features both sections of the SCAF10 IC configured as identical band-pass filters. The bandwidth (Q) is adjustable via R7A and R7B, while the center frequency is adjustable by R10 and the trimmer R9 as these resistors set the frequency of the LM555 clock generator.

As in the previous SCAF filter example, this filter also includes an audio amplifier (LM386) for driving either an external speaker or headphones as none of the SCAF filter ICs are capable of driving headphones directly.

10.4.2 Active RC Filters

Active RC filters based on op amp circuits can be used in either high-level audio output or very low-level direct-conversion receiver front-end filtering applications and thus are extremely flexible. Unlike LC filters, active RC filters can provide both filtering and gain at the same time, eliminating the relatively high insertion loss of a physically small, sharp LC filter. In addition, an active RC filter is not susceptible to the same ac hum pickup in low signal level applications as LC filters.

Active filters can be designed for gain and they offer excellent stage-to-stage isolation. The circuits require only resistors and capacitors, avoiding the limitations associated with inductors. By using gain and feedback, filter Q is controllable to a degree unavailable to passive LC filters. Despite the advantages, there are also some limitations. They require power, and performance may be limited by the op amp's finite input and output levels, gain and bandwidth. Active filters that drive speakers or other heavy loads usually employ an audio output amplifier circuit to boost the output power after a low-power filter stage.

A particular advantage of active RC filters for use in receivers is that they can be capable of extremely low noise operation, allowing the filtering of extremely small signals. At the same time they are capable of handling extremely large signals. Op amps are often capable of using ±18-V dual supply voltages or 36-V single supply voltages allowing the construction of a very high performance audio filter/amplifier chain that can handle signals up to 33 $V_{P\text{-}P}$. The ability to provide gain while also providing filtering provides a lot of flexibility in managing the sensitivity of a receiver audio chain.

The main disadvantage of active RC filters is their relatively high parts count compared to other filter types. In addition, active RC filters tend to be fixed-frequency designs unlike SCAF filters whose frequency can be moved simply by changing the clock frequency that drives the SCAF IC.

10.4.3 Active Filter Responses

Active filters can implement any of the passive LC filter responses described in the preceding section: low-pass, high-pass, band-

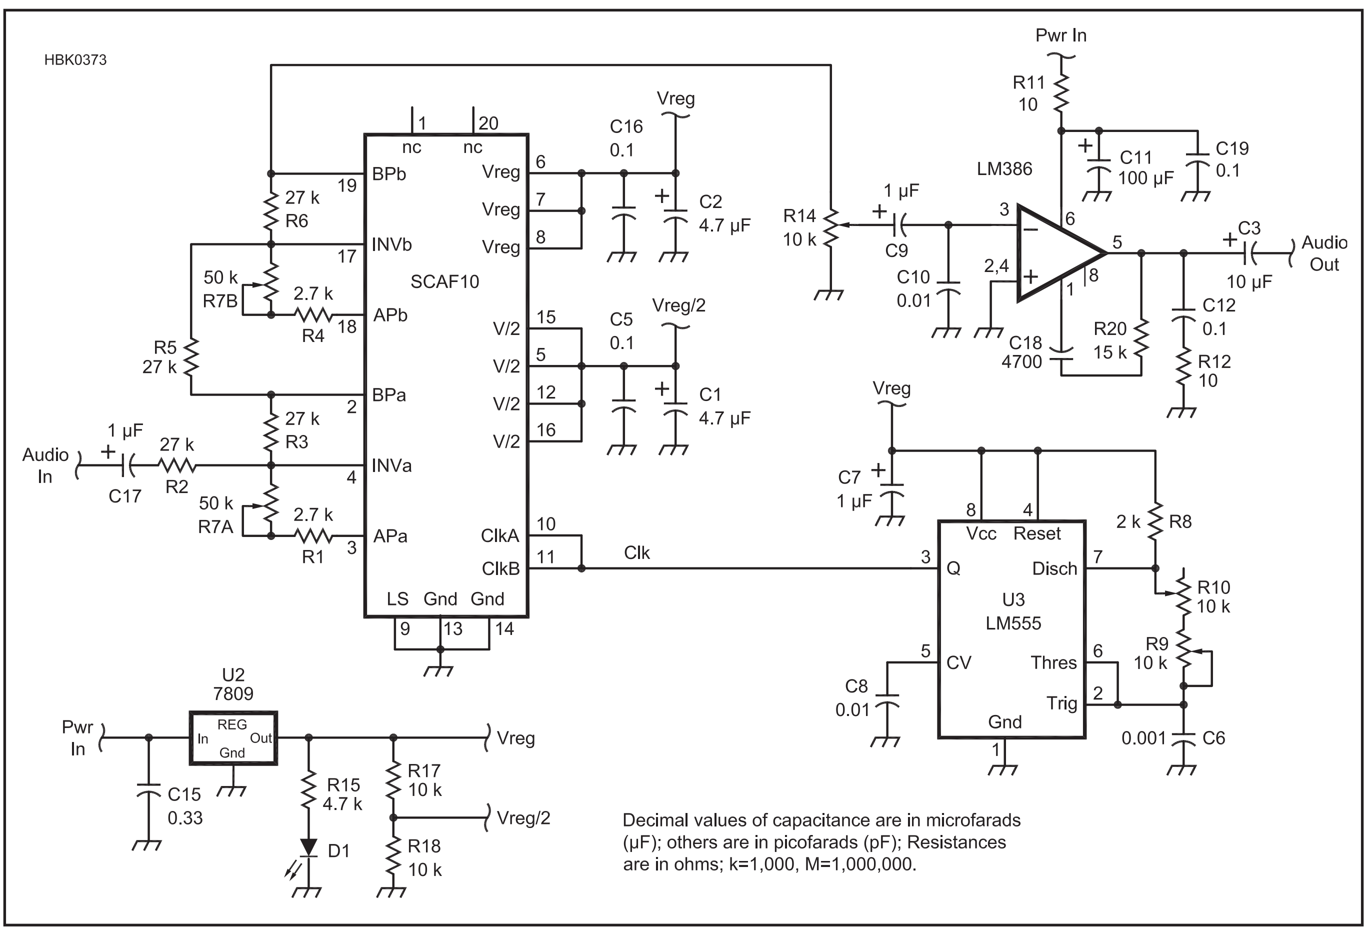

Figure 10.30 — Example of a simple SCAF band-pass filter from a New England QRP Club kit.

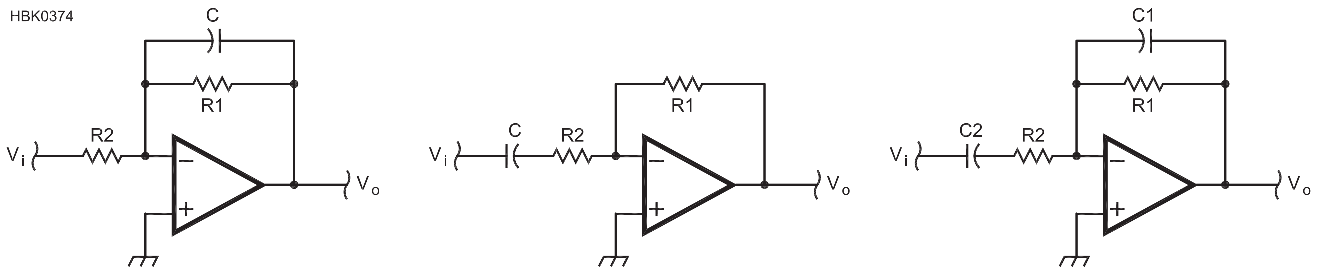

Passband gain, $A_v = -R1/R2$
Cutoff Frequency = $1/2\pi R1C$
(A)

Passband gain, $A_v = -R1/R2$
Cutoff Frequency = $1/2\pi R2C$
(B)

Passband gain, $A_v = -R1/R2$
Bandwidth = $1/2\pi R2C - 1/2\pi R1C$
(C)

Figure 10.31 — Simple active filters. A low-pass filter is shown at A; B is a high-pass filter. The circuit at C combines the low- and high-pass filters into a wide band-pass filter.

Unless otherwise specified, values of R are in ohms, C is in farads, F in hertz and ω in radians per second. Calculations shown here were performed on a scientific calculator.

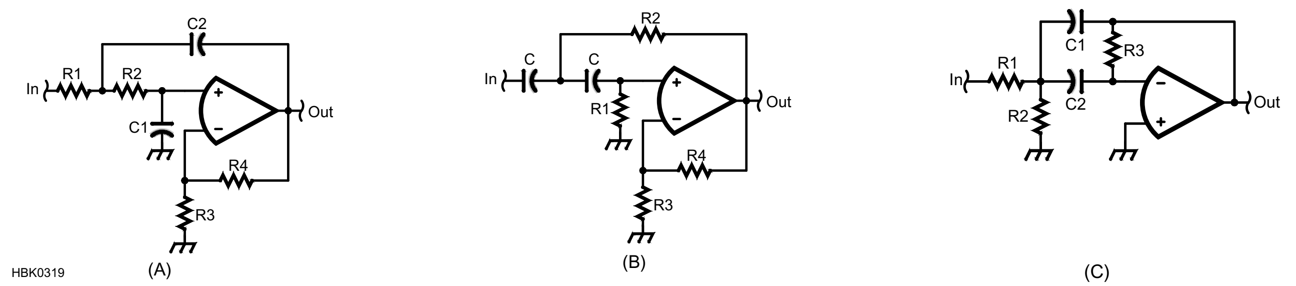

Low-Pass Filter

$$C_1 \le \frac{\left[a^2 + 4(K-1)\right]C_2}{4}$$

$$R_1 = \frac{2}{\left[aC_2 + \sqrt{\left[a^2 + 4(K-1)\right]C_2{}^2 - 4C_1C_2}\right]\omega_C}$$

$$R_2 = \frac{1}{C_1C_2R_1\omega_C{}^2}$$

$$R_3 = \frac{K(R_1 + R_2)}{K-1} \quad (K > 1)$$

$$R_4 = K(R_1 + R_2)$$

where
- K = gain
- f_c = –3 dB cutoff frequency
- $\omega_c = 2\pi f_c$
- C_2 = a standard value near $10/f_c$ (in µF)

Note: For unity gain, short R4 and omit R3.

Example:
- a = 1.414 (see table, one stage)
- K = 2
- f = 2700 Hz
- ω_c = 16,964.6 rad/sec
- C_2 = 0.0033 µF
- C1 ≤ 0.00495 µF (use 0.0050 µF)
- R1 ≤ 25,265.2 Ω (use 24 kΩ)
- R2 = 8,420.1 Ω (use 8.2 kΩ)
- R3 = 67,370.6 Ω (use 68 kΩ)
- R4 = 67,370.6 Ω (use 68 kΩ)

High-Pass Filter

$$R_1 = \frac{4}{\left[a + \sqrt{a^2 + 8(K-1)}\right]\omega_C\, C}$$

$$R_2 = \frac{1}{\omega_C{}^2C^2R_1}$$

$$R_3 = \frac{KR_1}{K-1} \quad (K > 1)$$

$$R_4 = KR_1$$

where
- K = gain
- f_c = –3 dB cutoff frequency
- $\omega_c = 2\pi f_c$
- C = a standard value near $10/f_c$ (in µF)

Note: For unity gain, short R4 and omit R3.

Example:
- a = 0.765 (see table, first of two stages)
- K = 4
- f = 250 Hz
- ω_c = 1570.8 rad/sec
- C = 0.04 µF (use 0.039 µF)
- R1 = 11,123.2 Ω (use 11 kΩ)
- R2 = 22,722 Ω (use 22 kΩ)
- R3 = 14,830.9 Ω (use 15 kΩ)
- R4 = 44,492.8 Ω (use 47 kΩ)

Band-Pass Filter

Pick K, Q, $\omega_o = 2\pi f_c$
where f_c = center freq.
Choose C
Then

$$R1 = \frac{Q}{K_0\omega_0 C}$$

$$R2 = \frac{Q}{\left(2Q^2 - K_0\right)\omega_0 C}$$

$$R3 = \frac{2Q}{\omega_0 C}$$

Example:
- K = 2, f_o = 800 Hz, Q = 5 and C = 0.022 µF
- R1 = 22.6 kΩ (use 22 kΩ)
- R2 = 942 Ω (use 910 Ω)
- R3 = 90.4 kΩ (use 91 kΩ)

Figure 10.32 — Equations for designing a low-pass RC active audio filter are given at A. B, C and D show design information for high-pass, band-pass and band-reject filters, respectively. All of these filters will exhibit a Butterworth response. Values of K and Q should be less than 10. See Table 10.3 for values of "a."

pass, band-stop and all-pass. Filter family responses such as Butterworth, Chebyshev, Bessel and Cauer (elliptic) can be realized. All of the same family characteristics apply equally to passive and active filters and will not be repeated here. (Op amps are discussed in the **Circuits and Components** chapter.)

Figure 10.31 presents circuits for first-order low-pass (Figure 10.31A) and high-pass (Figure 10.31B) active filters. The frequency response of these two circuits is the same as a parallel and series RC circuit, respectively, except that these two circuits can have a pass-band gain greater than unity. Rolloff is shallow at 6 dB/octave. The two responses can be combined to form a simple band-pass filter (Figure 10.31C). This combination cannot produce sharp band-pass filters because of the shallow rolloff.

To achieve high-order responses with steeper rolloff and narrower bandwidths, more complex circuits are required in which combinations of capacitors and resistors create *poles* and *zeroes* in the frequency response. (Poles and zeros are described in the **Radio Fundamentals** chapter.) The various filter response families are created by different combinations of additional poles and zeroes. There are a variety of circuits that can be configured to implement the equations that describe the various families of filter responses. The most common circuits are *Sallen-Key* and *multiple-feedback*, but there are numerous other choices.

There are many types of active filters — this section presents some commonly used circuits as examples. A book on filter design (see the References section) will present more choices and how to develop designs based on the different circuit and filter family types. In addition, op amp manufacturer's publish numerous application notes and tutorials on active filter design. Several are listed in the References section of this chapter. The set of tutorials published by Analog Devices is particularly good.

SECOND-ORDER ACTIVE FILTERS

Figure 10.32 shows circuits for four second-order filters: low-pass (Figure 10.32A), high-pass (Figure 10.32B), band-pass (Figure 10.32C) and band-reject or notch (Figure 10.32D). Sequences of these filters are used to create higher-order circuits by connecting them in series. Two second-order filter stages create a fourth-order filter, and so forth.

The low-pass and high-pass filters use the Sallen-Key circuit. Note that the high-pass circuit is just the low-pass circuit with the positions of R1-R2 and C1-C2 exchanged. R3 and R4 are used to control gain in the low- and high-pass configurations. The band-pass filter is a multiple-feedback design. The notch filter is based on the twin-T circuit. All of the filter design equations and tables will result in a Butterworth family response. (For the Chebyshev and other filter responses, consult the references listed at the end of the chapter.)

The circuits in Figure 10.32 assume the op amp is operating from a balanced, bipolar power supply, such as ±12 V. If a single supply is used (such as +12 V and ground), the circuit must have a dc offset added and blocking capacitors between filter sections to prevent the dc offset from causing the op amp to saturate. The references listed at the end of the chapter provide more detailed information on single-supply circuit design.

Avoid electrolytic or tantalum capacitors as frequency-determining components in active filter design. These capacitors are best used for bypassing and power filtering as their tolerance is generally quite low, they have significant parasitic effects, and are usually polarized. Very small values of capacitance (less than 100 pF) can be affected by stray capacitance to other circuit components and wiring. High-order and high-Q filters require close attention to component tolerance and temperature coefficients, as well.

SECOND-ORDER ACTIVE FILTER DESIGN PROCEDURES

The following simple procedures are used to design filters based on the schematics in Figure 10.32. Equations and a design example are provided in the figure.

Low- and High-Pass Filter Design

To design a low- or high-pass filter using the circuits in Figure 10.32, start by determining your performance requirements for filter order (2, 4, 6, or 8), gain (K) and cutoff frequency (f_C). Calculate $\omega_C = 2\pi f_C$ and C_2 or C as required.

Table 10.3 provides design coefficients to create the Butterworth response from successive Sallen-Key low- and high-pass stages. A different coefficient is used for each stage. Obtain design coefficient "a" from Table 10.3. Calculate the remaining component values from the equations provided.

Band-Pass Filter Design

To design a band-pass filter as shown in Figure 10.32C, begin by determining the filter's required Q, gain, and center frequency, f_0. Choose a value for C and solve for the resistor values. Very high or very low values of resistance (above 1 MΩ or lower than 10 Ω) should be avoided. Change the value

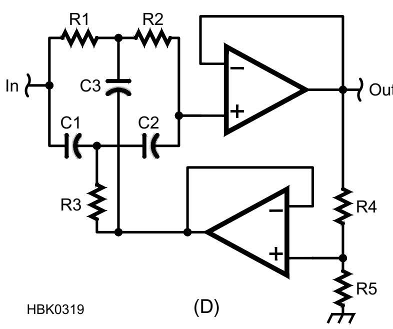

HBK0319 (D)

Band-Reject Filter

$$F_0 = \frac{1}{2\pi R1C1}$$

$$K = 1 - \frac{1}{4Q}$$

$$R >> (1-K)R1$$

where

$$C1 = C2 = \frac{C3}{2} = \frac{10\,\mu F}{f_0}$$

R1 = R2 = 2R3
R4 = (1 – K)R
R5 = K × R

Example:
f_0 = 500 Hz, Q=10
K = 0.975
C1 = C2 = 0.02 µF (or use 0.022 µF)
C3 = 0.04 µF (or use 0.044 µF)
R1 = R2 = 15.92 kΩ (use 15 kΩ)
R3 = 7.96 kΩ (use 8.2 kΩ)
R >> 1 kΩ
R4 = 25 Ω (use 24 Ω)
R5 = 975 Ω (use 910 Ω)

Table 10.3
Factor "a" for Low- and High-Pass Filters in Figure 10.32

No. of Stages	*Stage 1*	*Stage 2*	*Stage 3*	*Stage 4*
1	1.414	–	–	–
2	0.765	1.848	–	–
3	0.518	1.414	1.932	–
4	0.390	1.111	1.663	1.962

These values are truncated from those of Appendix C of Williams, *Electronic Filter Design Handbook*, for even-order Butterworth filters of C until suitable values are obtained. High gain and Q may be difficult to obtain in the same stage with reasonable component values. Consider a separate stage for additional gain or to narrow the filter bandwidth.

Band-Reject (Notch) Filter Design

Band-reject filter design begins with the selection of center frequency, f_0, and Q. Calculate the value of K. Choose a value for C1 that is approximately 10 µF / f_0. This determines the values of C2 and C3 as shown. Calculate the value for R1 that results in the desired value for f_0. This determines the values of R2 and R3 as shown. Select a convenient value for R4 + R5 that does not load the output amplifier. Calculate R4 and R5 from the value of K.

The depth of the notch depends on how closely the values of the components match the design values. The use of 1%-tolerance resistors is recommended and, if possible, matched values of capacitance. If all identical components are used, two capacitors can be paralleled to create C3 and two resistors in parallel create R3. This helps to minimize thermal drift.

10.4.4 Active Filter Design Tools

While the simple active filter examples presented above can be designed manually, more sophisticated circuits are more easily designed using filter-design software. Follow the same general approach to determining the filter's performance requirements and then the filter family as was presented in the section on LC filters. You can then enter the values or make the necessary selections for the design software. Once a basic design has been calculated, you can then "tweak" the design performance, use standard value components and make other adjustments. The design example presented below shows how a real analog design is assembled by understanding the performance requirements and then using design software to experiment for a "best" configuration.

Op amp manufacturers such as Texas Instruments and National Semiconductor (originally separate companies but now merged) have made available sophisticated "freeware" filter design software. These packages are extremely useful in designing active RC filters. They begin by collecting specifications from the user and then creating a basic circuit. Once the initial circuit has been designed, the user can adjust specifications, component values, and op amp types until satisfied with the final design.

Free online filter design software is available from the following sources:

• Analog Devices — *Analog Filter Wizard 2.0* (**tools.analog.com/en/filterwizard/**)

• Texas Instruments — *FilterPro 3.1* (**www.ti.com**, search for "FilterPro")

Even if a manual design process is followed, using a software tool to double-check the results is a good way to verify the design before building the circuit. A circuit simulator such as *LTSpice* (see the **Computer-Aided Circuit Design** chapter) can also verify a design before building the actual circuit.

An extended design example using the Texas Instruments *FilterPro* software is included in this book's online information. In the example, Dan Tayloe, N7VE explains the process of designing a high-performance 750 Hz low-pass filter, illustrating the power of using sophisticated interactive tools that enable design changes on-the-fly. The reader is encouraged to follow along and experiment with *FilterPro* as a means of becoming familiar with the software so that it can be used for other filter design tasks. Similar processes apply to other filter design software tools.

10.5 Digital Filters

Where an analog filter operates on a continuous signal in the time domain, digital filters operate on signals that have been converted to a digital stream of data. (See the discussion of digital signals in the **DSP and SDR Fundamentals** chapter.) The usual method of constructing a digital filters is as a series of registers with the stream of samples moving through them in a series of steps called *delays*. At each register (called a *tap*) the amplitude of the signal is modified (by an amount referred to as a *coefficient*) and the result added to the results from other taps. The signal stream is shifted to the next register and the process is repeated. The sum of the results from the taps constitutes the output of the filter.

The value of each coefficient and the configuration of how the filter's taps are added together determine the filter's response to the input signal. This section will discuss two basic filter types, the finite impulse response or FIR, the infinite impulse response or IIR.

The purpose of this section is not to give a tutorial in how each type of filter is designed but explain the characteristics of each type of filter. As with analog filters, the design of digital filter is mostly done by software using validated algorithms that accept performance requirements as inputs and generate design information, in many cases as a data file that can be loaded directly into a DSP processor or subsystem. (The GNU Radio (**gnuradio.org**) toolkit, for example, treats digital filters as one of many functional blocks. The user configures the filter parameters and the software manages the organization of the actual computing elements in the DSP or SDR system.) For the reader interested in a deeper discussion of digital filter theory and design, several excellent references are listed at the end of the chapter.

10.5.1 FIR Filters

The earlier discussion of transient response introduced the impulse, an infinitely narrow pulse. The spectrum of an impulse is infinitely wide, containing all frequencies. If you feed an impulse into the input of a filter (analog or digital), the output of the filter is the filter's impulse response. Because the impulse is made up of all frequencies, it excites the filter at all frequencies. As a result, the impulse response has a frequency spectrum equal to the frequency response of the filter.

One way to design a filter is to determine the impulse response that corresponds to the desired frequency spectrum and then design the filter to have that impulse response. That method is the basis for designing *finite impulse response* (FIR) filters. (Finite refers to time, not amplitude.)

An FIR filter is a filter whose impulse response is finite, ending in some fixed time. Note that analog filters have an infinite impulse response — the output theoretically rings forever. Even a simple R-C low-pass filter's output decays exponentially toward zero but theoretically never quite reaches it. In contrast, an FIR filter's impulse response becomes exactly zero at some time after receiving the impulse and stays zero forever (or at least until another impulse comes along).

Given that you have somehow figured out the desired impulse response, how would you design a digital filter to have that response? The obvious method would be to pre-calculate a table of impulse response values, sampled at the sample rate. These are called the filter *coefficients*. When an impulse of a certain amplitude is received, you multiply that amplitude by the first entry in the coefficient table and send the result to the output. At the next sample time, multiply the impulse by the second entry, and so on until you have used up all the entries in the table.

A circuit to do that is shown in **Figure 10.33**. The input signal is stored in a shift register. Each block labeled "Delay" represents a delay of one sample time. At each sample time, the signal is shifted one register to the right. Each register feeds a multiplier and the other input to the multiplier comes from one of the coefficient table entries. All

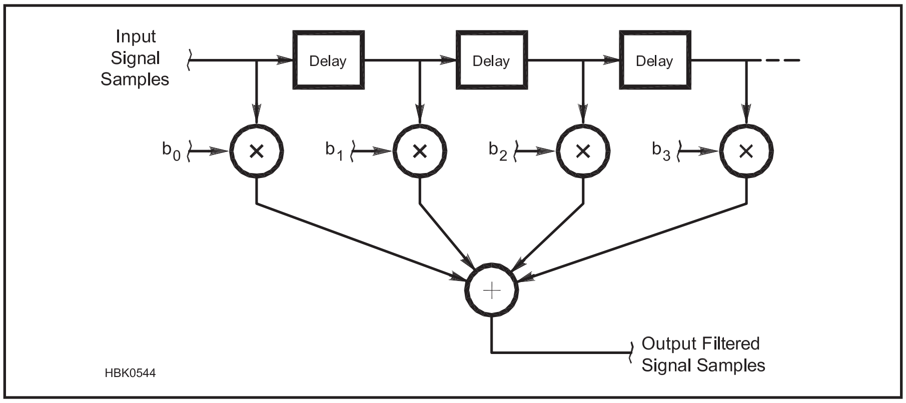

Figure 10.33 — A 4-tap FIR filter. The bn values are the filter coefficients.

the multiplier outputs are added together. Since the input is assumed to be a single impulse, at any given time all the shift registers contain zero except one, which is multiplied by the appropriate table entry and sent to the output.

We've just designed an FIR filter! By using a shift register with a separate multiplier for each tap, the filter works for continuous signals as well as for impulses. Since this is a linear system, the continuing signal is affected by the filter the same as an individual impulse.

It should be obvious from the diagram how to implement an FIR filter in software. You set up two buffers in memory, one for the filter coefficients and one for the data. The length of each buffer is the number of filter taps. (A *tap* is the combination of one filter coefficient, one shift register and one multiplier/accumulator.) Each time a new data value is received, it is stored in the next available position in the data buffer and the accumulator is set to zero. Next, a software loop is executed a number of times equal to the number of taps. During each loop, pointers to the two buffers are incremented, the next coefficient is multiplied by the next data value and the result is added to the current accumulator value. After the last loop, the accumulator contents are the output value. Normally the buffers are implemented as *circular buffers* — when the address pointer gets to the end it is reset back to the beginning.

Now you can see why a hardware multiplier-accumulator (MAC) is such an important feature of a DSP chip. Each tap of the FIR filter involves one multiplication and one addition. With a 1000-tap FIR filter, 1000 multiplications and 1000 additions must be performed during each sample time. The ability to do each MAC operation in a single clock cycle saves a lot of processing time compared to a CPU-type processor.

An FIR filter is a hardware or software implementation of the mathematical operation called *convolution*. We say that the filter *convolves* the input signal with the impulse response of the filter. It turns out that convolution in the time domain is mathematically equivalent to multiplication in the frequency domain. That means that the frequency spectrum of the output equals the frequency spectrum of the input times the frequency spectrum of the filter. Expressed in decibels, the output spectrum equals the input spectrum plus the filter frequency response, all in dB. If at some frequency the input signal is +3 dB and the filter is –10 dB compared to some reference, then the output signal will be 3 – 10 = –7 dB at that frequency.

An FIR filter whose bandwidth is very small compared to the sample rate requires a long impulse response with lot of taps. This is another consequence of the "skinny" versus "fat" relationship between the frequency and time domains. If the filter is narrow in the frequency domain, then its impulse response is wide. Actually, if you want the frequency response to go all the way to zero (minus infinity dB) throughout the stopband, then the impulse response theoretically becomes infinitely wide. Since we're designing a *finite* impulse response filter we have to truncate the impulse response at some point to get it to fit in the coefficient table. When you do that, however, you no longer have infinite attenuation in the stopband. The more heavily you truncate (the narrower the impulse response) the worse the stopband attenuation and the more ripple you get in the passband. Assuming optimum design techniques for selecting coefficients, you can estimate the minimum length L of the impulse response from the following equation:

$$L = 1 - \frac{10\log(\delta_1\delta_2) - 15}{14\left(\frac{f_T}{f_s}\right)} \text{ taps}$$

where

δ_1 and δ_2 = the passband and stopband ripple expressed as a fraction,

f_T = the transition bandwidth (frequency difference between passband and stopband edges), and

f_s = the sample rate.

For example, for a low-pass filter with a passband that extends up to 3 kHz, a stopband that starts at 4 kHz (f_T = 4 – 3 = 1 kHz), f_s = 10 kHz sample rate, ±0.1 dB passband ripple ($d_1 = 10^{0.1/20} - 1 = 0.0116$), and 60 dB stopband rejection ($d_2 = 10^{-60/20} = 0.001$), we get

$$L = 1 - \frac{10\log(0.0116 \times 0.001) - 15}{14\left(\frac{1}{10}\right)}$$

$$= 1 - \frac{-49.4 - 15}{1.4} = 47 \text{ taps}$$

Overflow is a potential problem when doing the calculations for an FIR filter. Multiplying two N-bit numbers results in a product with 2N bits, so space must be provided in the accumulator to accommodate that. Although the final result normally will be scaled and truncated back to N bits, it is best to carry through all the intermediate results with full resolution in order not to lose any dynamic range. In addition, the sum of all the taps can be a number with more than 2N bits. For example, if the filter width is 256 taps, then if all coefficients and data are at full scale, the final result could theoretically be 256 times larger, requiring an extra 8 bits in the accumulator. We say "theoretically" because normally most of the filter coefficients are much less than full scale and it is highly unlikely that all 256 data values would ever simultaneously be full-scale values of the correct polarity to cause overflow. The dsPIC processors use 16-bit multipliers with 32-bit results and a 40-bit accumulator, which should handle any reasonable circumstances.

After all taps have been calculated, the final result must be retrieved from the accumulator. Since the accumulator has much more resolution than the processor's data words, normally the result is truncated and scaled to fit. It is up to the circuit designer or programmer to scale by the correct value to avoid overflow. The worst case is when each data value in the shift register is full-scale — positive when it is multiplying a positive coefficient and negative for negative coefficients. That way, all taps add to the maximum value. To calculate the worst-case accumulator amplitude, simply add the absolute values of all the coefficients. However, that normally gives an unrealistically pessimistic value because statistically it is extremely unlikely that such a high peak will ever be reached. For a low-pass filter, a better estimate is to calculate the gain for a dc signal and add a few percent safety margin. The dc gain is just the sum of all the coefficients (not the absolute values). For a band-pass filter, add the sum of all the coefficients multiplied by a sine wave at the center frequency.

CALCULATING FIR FILTER COEFFICIENTS

So far we have ignored the question of how to determine the filter coefficients. For an ideal "brick-wall" low-pass filter, the answer turns out to be pretty simple. A "brick-wall" low-pass filter is one that has a constant response from zero hertz up to the cutoff frequency and zero response above. Its impulse response is proportional to the sinc function:

$$C(n) = C_o \ \text{sinc}(2Bn) = \frac{\sin(2\pi Bn)}{2\pi Bn}$$

where C(n) are the filter coefficients, n is the sample number with n = 0 at the center of the impulse response, C_o is a constant, and B is the single-sided bandwidth normalized to the sample rate, B = bandwidth / sample rate.

It is interesting that this has the same form as the frequency response of a pulse, as was shown in Figure 10.33. That is because a brick-wall response in the frequency domain has the same shape as a pulse in the time domain. A pulse in one domain transforms to a sinc function in the other. This is an example of the general principle that the transformation between time and frequency domains is symmetrical. This is discussed in more detail in the section on Fourier transforms in the **DSP and SDR Fundamentals** chapter.

Normally, the filter coefficients are set up with the peak of the sinc function, sinc(0), at the center of the coefficient table so that there is an equal amount of "tail" on both sides. That points up the principle problem with this method of determining filter coefficients. Theoretically, the sinc function extends from minus infinity to plus infinity. Abruptly terminating the tails causes the frequency response to differ from an ideal brick-wall filter. There is ripple in the passband and non-zero response in the stopband, as shown in the graph in the upper right of **Figure 10.34**. This is mainly caused by the abruptness of the truncation. In effect, all the coefficients outside the limits of the coefficient table have been set to zero. The passband and stopband response can be improved by tapering the edges of the impulse response instead of abruptly transitioning to zero.

The process of tapering the edges of the impulse response is called *windowing*. The impulse response is multiplied by a window, a series of coefficients that smoothly taper to zero at the edges. For example, a rectangular window is equivalent to no window at all. Many different window shapes have been developed over the years — at one time it seemed that every doctoral candidate in the field of signal processing did their dissertation on some new window. Each window has its advantages and disadvantages. A window that transitions slowly and smoothly to zero has excellent passband and stopband response but a wide transition band. A window that has a wider center portion and then transitions more abruptly to zero at the edges has a narrower transition band but poorer passband and stopband response. The equations for the windows in Figure 10.34 are included in a sidebar.

The routine shown in **Table 10.4** is written for a dsPIC processor so it can be used to calculate filter coefficients "on the fly" as the operator adjusts a bandwidth control. The same code should also work using a generic C compiler on a PC so the coefficients could be downloaded into an FIR filter implemented in hardware.

The windowed-sinc method works pretty well for a simple low-pass filter, but what if some more-complicated spectral shape is

Table 10.4
Routine for dsPIC Processor to Calculate Filter Coefficients

```
// Calculate FIR filter coefficients
// using the windowed-sinc method
void set_coef (
 double sample_rate;
 double bandwidth;)
{
extern int c[FIR_LEN]; // Coefficient array
int i; // Coefficient index
double ph; // Phase in radians
double coef; // Filter coefficient
int coef_int; // Digitized coefficient
double bw_ratio; // Normalized bandwidth

bw_ratio = 2 * bandwidth / sample_rate;
for (i = 0; i < (FIR_LEN/2); i++) {
 // Brick-wall filter:
 ph = PI * (i + 0.5) * bw_ratio;
 coef = sin(ph) / ph;
 // Hann window:
 ph = PI * (i + 0.5) / (FIR_LEN/2);
 coef *= (1 + cos(ph)) / 2;
 // Convert from floating point to int:
 coef *= 1 << (COEF_WIDTH - 1);
 coef_int = (int)coef;
 // Symmetrical impulse response:
 c[i + FIR_LEN/2] = coef_int;
 c[FIR_LEN/2 - 1 - i] = coef_int;
 }
}
```

Equations for Window Functions

For each window function, the center of the response is considered to be at time t = 0 and the width of the impulse is L. Each window is 1.0 when t = 0 and 0.0 when the | t | > L/2.

Rectangular:

w(t) = 1.0

Triangular (Bartlett):

$$w(t) = 2\left(\frac{L/2 - |t|}{L}\right)$$

Blackman:

$$w(t) = 0.42 + 0.5\cos\left(\frac{2\pi t}{L}\right) + 0.08\cos\left(\frac{4\pi t}{L}\right)$$

Hamming:

$$w(t) = 0.54 + 0.46\cos\left(\frac{2\pi t}{L}\right)$$

Hanning (Hann):

$$w(t) = 0.5 + 0.5\cos\left(\frac{2\pi t}{L}\right)$$

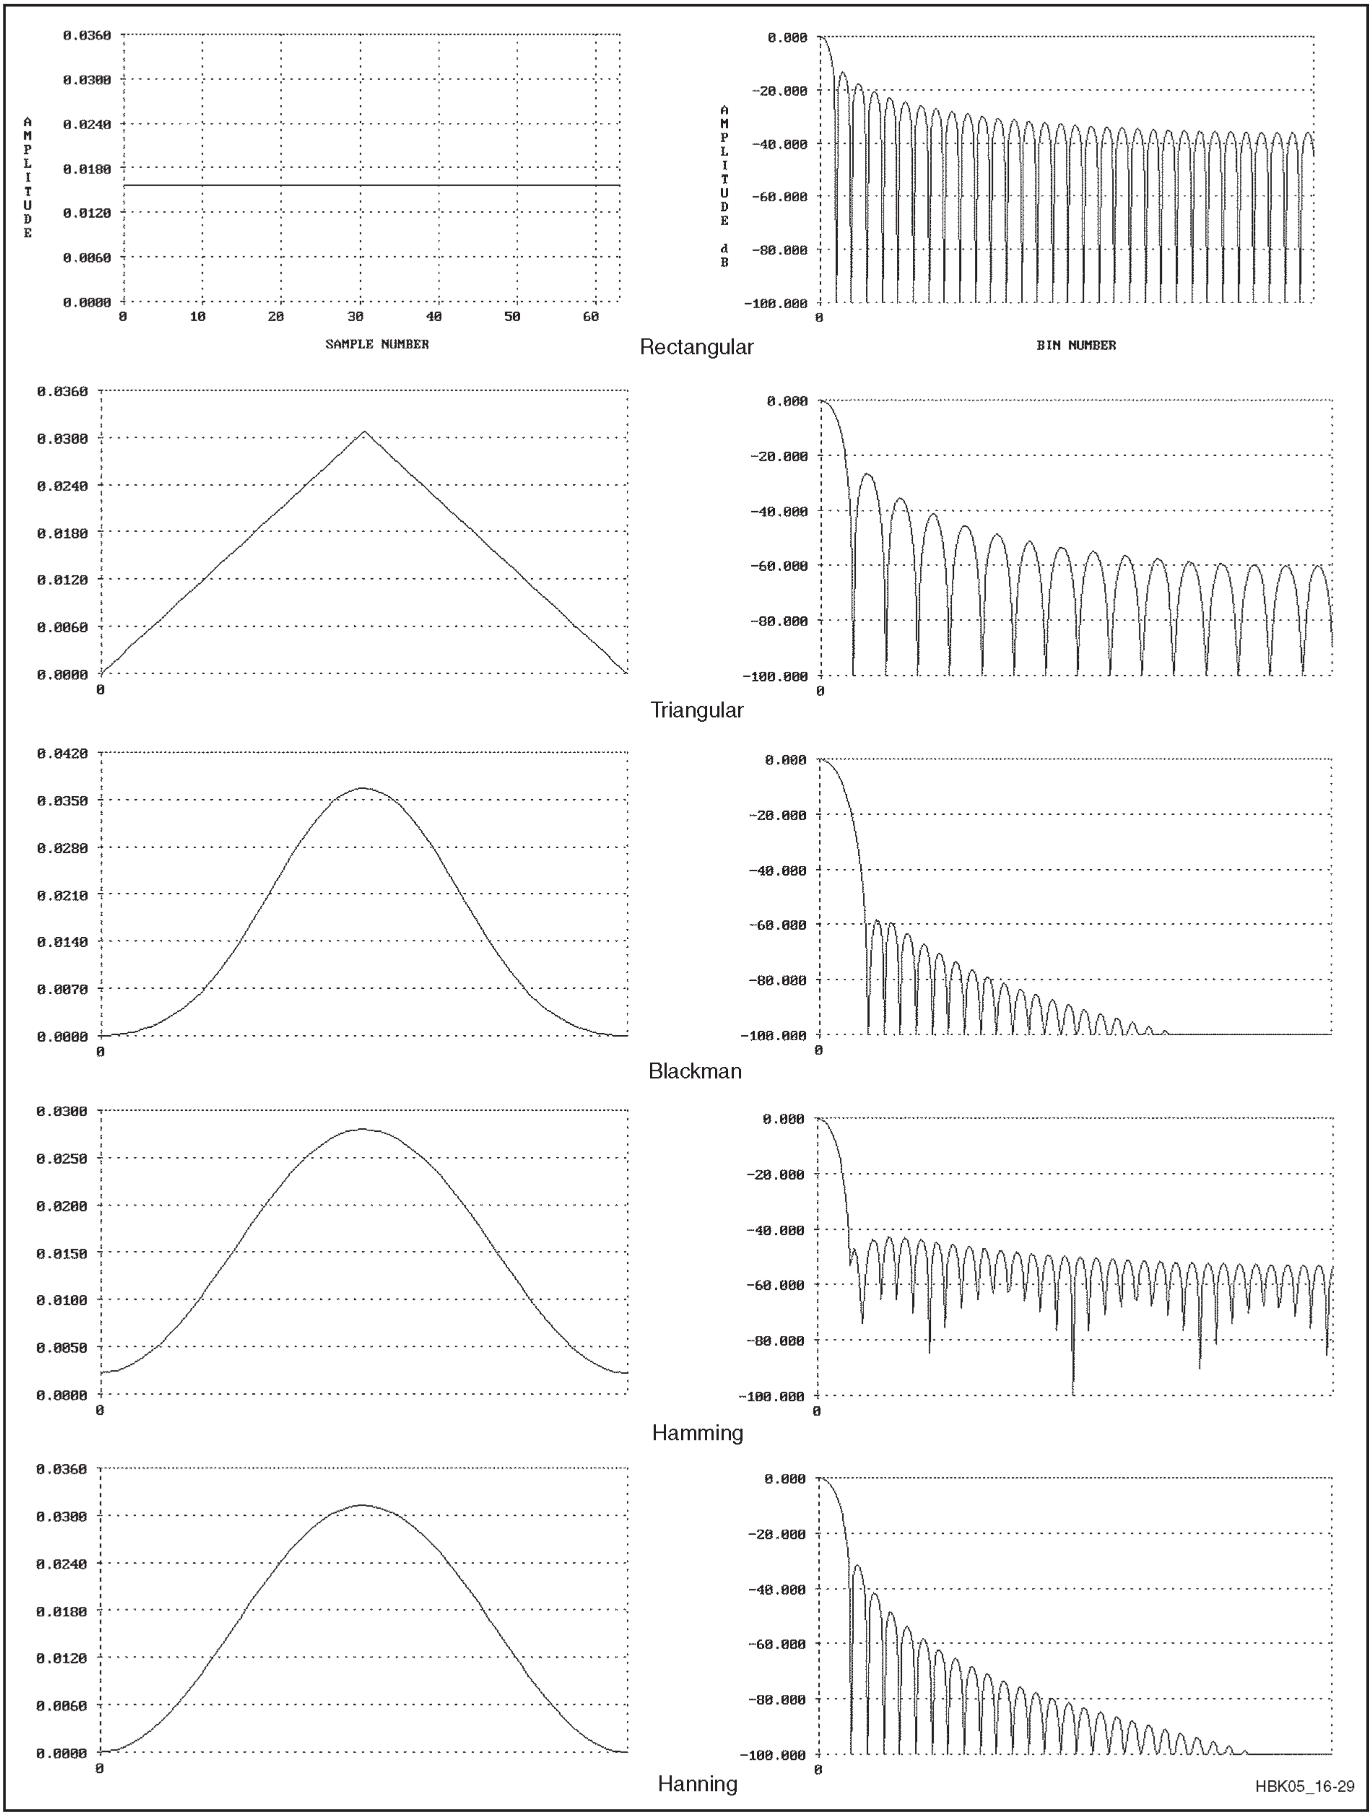

Figure 10.34 — Various window functions and their Fourier transforms.

desired? The same general design approach still applies. You determine the desired spectral shape, transform it to the time domain using an inverse Fourier transform, and apply a window. There is lots of (often free) software available that can calculate the fast Fourier transform (FFT) and inverse fast Fourier transform (IFFT). So the technique is to generate the desired spectral shape, transform to the time domain with the IFFT, and multiply the resulting impulse response by the desired window. Then you can transform back to the frequency domain with the FFT to see how the window affected the result. If the result is not satisfactory, you can either choose a different window or modify the original spectral shape and go through the process again.

Windowing methods are useful because they are simple to program and the resulting software routines execute quickly. For example, you can include a bandwidth knob on your DSP filter and calculate filter coefficients "on the fly" as the user turns the knob. However, while the filter performance that results is pretty good, it is not "optimum" in the sense that it does not have minimum passband ripple and maximum stopband attenuation for a given number of filter coefficients. For that, you need what is known as an *equal-ripple*, or *Chebyshev* filter. The calculations to determine Chebyshev filter coefficients are more complicated and time-consuming. For that reason, the coefficients are normally calculated in advance on a PC and stored in DSP program memory for retrieval as needed.

Engineers have not had much success in devising a mathematical algorithm to calculate the Chebyshev coefficients directly, but in 1972 Thomas Parks and James McClellan figured out a method to do it iteratively. The *Parks-McClellan algorithm* is supported in most modern filter-design software, including a number of programs available for free download on the Web. Typically you enter the sample rate, the passband and stopband frequency ranges, the passband ripple and the stopband attenuation. The software then determines the required number of filter coefficients, calculates them and displays a plot of the resulting filter frequency response.

Filter design software typically presents the filter coefficients as floating-point numbers to the full accuracy of the computer. You will need to scale the values and truncate the resolution to the word size of your filter implementation. Truncation of filter coefficients affects the frequency response of the filter but does not add noise in the same manner as truncating the signal data.

As you look at impulse responses for various FIR filters calculated by various methods you soon realize that most of them are symmetrical. If the center of the impulse response is considered to be at time zero, then the value at time t equals the value at time –t for all t.

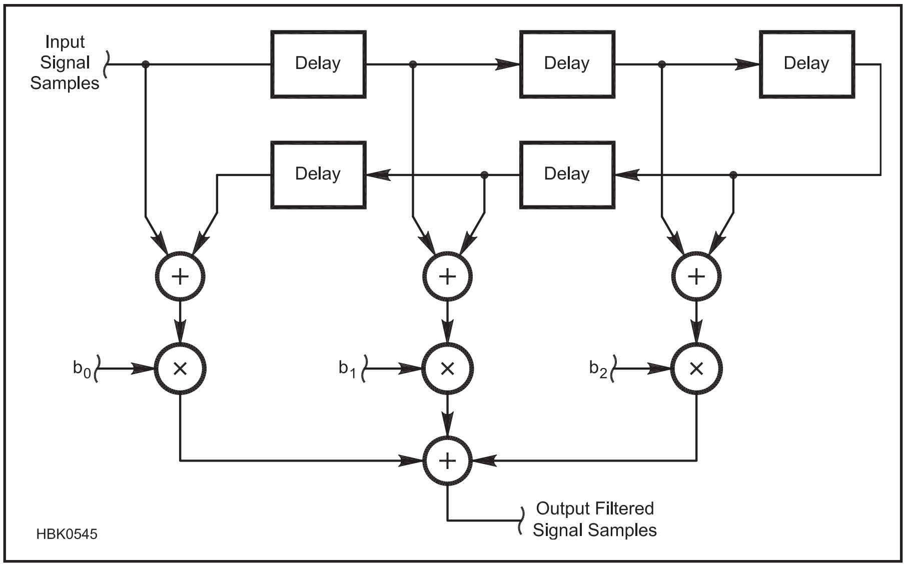

Figure 10.35 — A 6-tap FIR filter. Because the coefficients are symmetrical, the symmetrical taps may be combined before multiplication.

If you know in advance that the filter coefficients are symmetrical, you can take advantage of that in the filter design. By re-arranging the adders and multipliers, the number of multipliers can be reduced by a factor of two, as shown in **Figure 10.35**. This trick is less useful in a software implementation of an FIR filter because the number of additions is the same and many DSPs take the same amount of time to do an addition as a multiply-accumulate.

In addition to the computational benefit, a symmetrical impulse response also has the advantage that it is *linear phase*. The time delay through such a filter is one-half the length of the filter for all frequencies. For example, for a 1000-tap filter running at 10 kHz the delay is 500/10,000 = 0.05 second. Since the time delay is constant for all frequencies, the phase delay is directly proportional to the frequency. For example, if the phase delay at 20 Hz is one cycle (0.05 second) it is ten cycles at 200 Hz (still 0.05 second). Linear phase delay is important with digital modulation signals to avoid distortion and inter-symbol interference. It is also desirable with analog modulation where it can result in more natural-sounding audio. All analog filters are non-linear-phase; the phase distortion tends to be worse the more abrupt the transition between passband and stopband. That is why an SSB signal sounds unnatural after being filtered by a crystal filter with a small shape factor even though the passband ripple may be small and distortion minimal.

A band-pass filter can be constructed from a low-pass filter simply by multiplying the impulse response by a sine wave at the desired center frequency. This can be done before or after windowing. The linear-phase property is retained but with reference to the center frequency of the filter, that is, the phase shift is proportional to the difference in frequency from the center frequency. The frequency response is a double-sided version of the low-pass response with the zero-hertz point of the low-pass filter shifted to the frequency of the sine wave.

10.5.2 IIR Filters

An *infinite impulse response* (IIR) filter is a filter whose impulse response is infinite. After an impulse is applied to the input, theoretically the output never goes to zero and stays there. In practice, of course, the signal eventually does decay until it is below the noise level (analog filter) or less than one LSB (digital filter).

Unlike a symmetrical FIR filter, an IIR filter is not generally linear-phase. The delay through the filter is not the same for all frequencies. Also, IIR filters tend to be harder to design than FIR filters. On the other hand, many fewer adders and multipliers are typically required to achieve the same passband and stopband ripple in a given filter, so IIR filters are often used where computations must be minimized.

All analog filters have an infinite impulse response. For a digital filter to be IIR it must have feedback. That means a delayed copy of some internal computation is applied to an earlier stage in the computation. A simple but useful example of an IIR filter is the exponential decay circuit in **Figure 10.36**. In the absence of a signal at the input, the output on the next clock cycle is always (1–δ) times the current output. The time constant (the time for the output to die to 1/e = 36.8% of the initial value) is very nearly

$$\tau = f_s\left(\frac{1}{\delta} - \frac{1}{2}\right)$$

where f_s is the sample rate. The circuit is the digital equivalent of a capacitor with a resistor in parallel and might be useful for example in a digital automatic gain control circuit.

One issue with IIR filters is resolution. Because of the feedback, the number of bits of resolution required for intermediate computations can be much greater than at the input or output. In the previous example, δ is very small for very long time constants. When the value in the register falls below a certain level the multiplication by (1–δ) will no longer be accurate unless the bit width is increased. In practice, the increased resolution required with IIR filters often cancels out part of the savings in the number of circuit elements.

Another issue with IIR filters is stability. Because of the feedback it is possible for the filter to oscillate if care is not taken in the design. Stability can also be affected by nonlinearity at low signal levels. A circuit that is stable with large signals may oscillate with small signals due to the round-off error in certain calculations, which causes faint tones to appear when strong signals are not present. This is known as an unstable *limit cycle*. These issues are part of the reason that IIR filters have a reputation for being hard to design.

Design techniques for IIR filters mostly involve first designing an analog filter using any of the standard techniques and then transforming the design from the analog to the digital domain. The *impulse-invariant* method attempts to duplicate the filter response directly by making the digital impulse response equal the impulse response of the equivalent analog filter. It works fairly well for low-pass filters with bandwidths much less than the sample rate. Its problem is that it tries to duplicate the frequency response all the way to infinity hertz, but that violates the Nyquist criterion resulting in a folding back of the high-frequency response down into low frequencies. It is similar to the aliasing that occurs in a DSP system when the input signal to be sampled is not band-limited below the Nyquist frequency.

The *bilinear transform* method gets around that problem by distorting the frequency axis such that infinity hertz in the analog domain becomes sample rate / 2 in the digital domain. Low frequencies are fairly accurate, but high frequencies are squeezed together more and more the closer you get to the Nyquist frequency. It avoids the aliasing problem at the expense of a change in the spectrum shape, especially at the high-frequency end. For example, when designing a low-pass filter it may be necessary to change the cutoff frequency to compensate. Again, the method works best for filters with passband frequencies much less than the sample rate.

In general, the output of an IIR filter is a combination of the current and previous input values (feed-forward) and previous output values (feed-back). **Figure 10.37A** shows the so-called *direct form I* of an IIR filter. The b_i coefficients represent feed-forward and the a_i coefficients feed-back. For example the previous value of the y output is multiplied by a_1, the second previous value is multiplied by a_2, and so on. Because the filter is linear, it doesn't matter whether the feed forward or feed back stage is performed first. By reversing the order, the number of shift registers is reduced as in Figure 10.37B. There are other equivalent topologies as well. The mathematics for generating the a_i and b_i coefficients for both the impulse-invariant and bilinear transform methods is fairly involved, but fortunately some filter design programs can handle IIR as well as FIR filters.

10.5.3 CIC Filters

(This overview of CIC filters is taken from the 2011 Sep/Oct *QEX* column SDR: Simplified by Ray Mack, W5IFS. Additional background on CIC filters is available in the referenced articles by Donadio and by Lyons.)

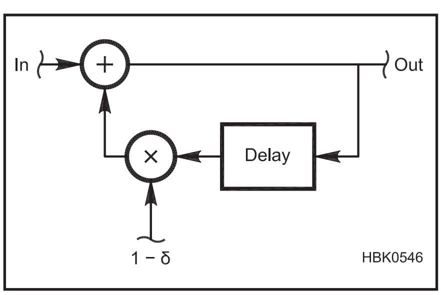

Figure 10.36 — An exponential decay implementation.

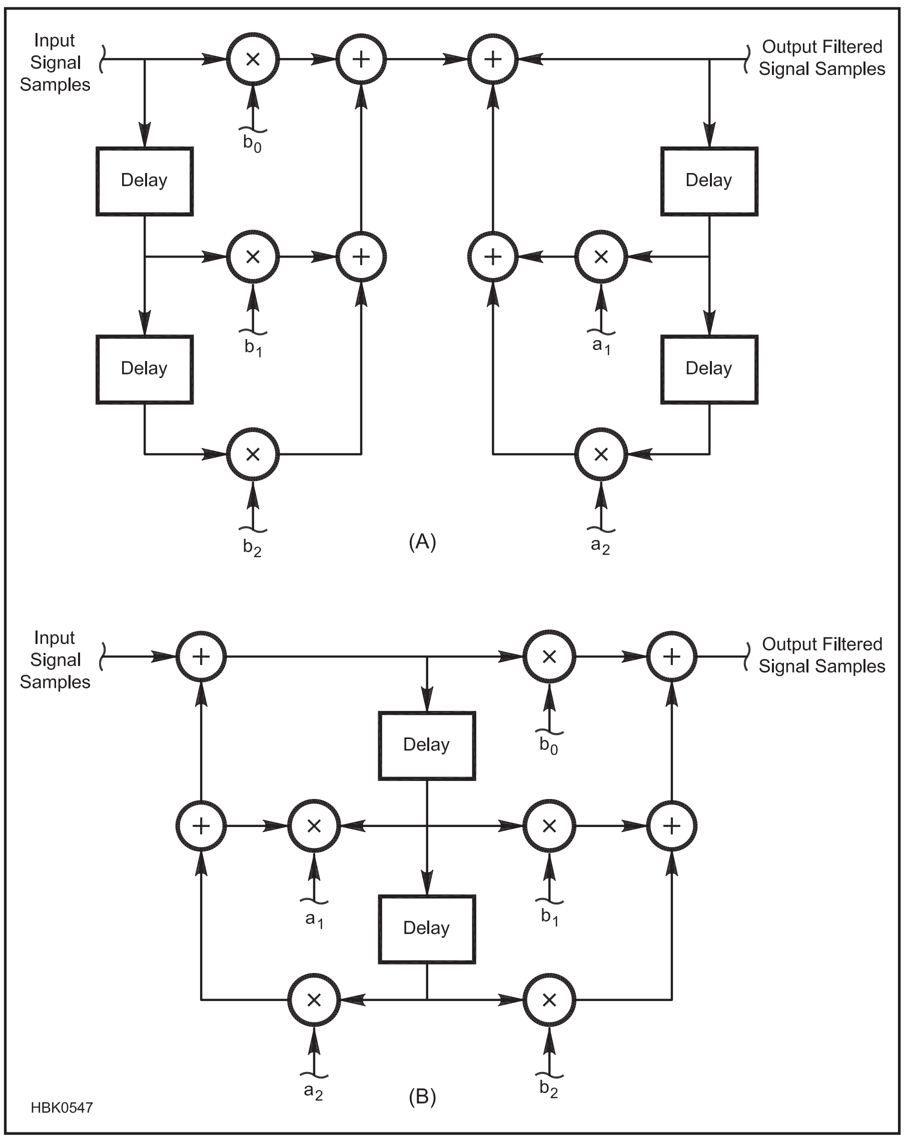

Figure 10.37 — An IIR filter with three feed-forward taps and two feed-back taps. Direct form I (A) and the equivalent direct form II (B).

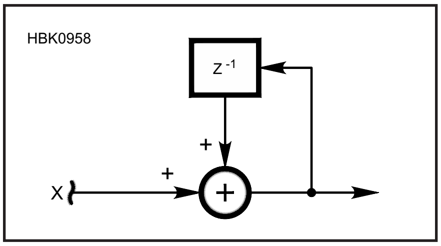

Figure 10.38 — Z transform diagram of an integrator. The new output is the sum of the previous output (represented by Z^{-1}) and the new sample. (The Z transform is explained at en.wikipedia.org/wiki/Z-transform.)

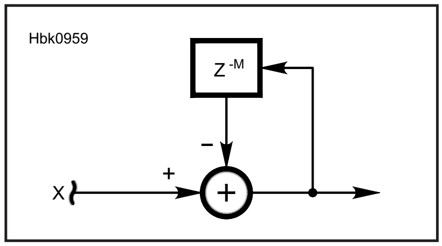

Figure 10.39 — Z transform diagram of a comb. The new output is the difference of the present sample and a delayed sample. The number M designates how many steps happened during the delay.

The problem with both FIR and IIR filters is that they require a large number of multiply operations, which consume a large number of DSP processor cycles. Eugene Hogenauer (see references) developed a very useful simplification of the sample conversion/filter configuration called a *cascaded integrator comb* (CIC) filter. The important aspect of CIC filters is that only addition, subtraction, and delay operations are required for implementation.

As with most things in life, improving one aspect of a system requires compromise in other aspects. This is also true of CIC filters where we trade the simplification of eliminating multipliers for restricting the filter response: A CIC filter can only have a low-pass response. Additionally, there is a limited subset of possible low-pass responses constrained by the sample rate change and number of stages in the comb and integrator stages. The most important property of a CIC filter is that it can be very easily implemented in hardware either in an FPGA or as part of the dedicated logic of an IC.

The integrator is an infinite impulse response filter. **Figure 10.38** shows how it works and how simple it is. The integrator holds a running total of all previous samples. The integrator adds the last output value (z–1) to the current input value (x). Ordinarily, we would worry about overflow in an integrator because a dc component in the signal will cause the integrator to overflow.

The combination of the comb and the integrator, however, cancels any problems with overflow. The integrator is a single pole low pass filter with infinite gain at dc. (The possibility of overflow doesn't matter as long as the integrators are implemented with adders using two's complement addition (**en.wikipedia.org/wiki/Two's_complement**) that allow wrap around when overflow occurs, and that the number of bits in the word is as big as the expected output word.)

The comb is a finite impulse response stage that subtracts a previous sample from the present sample. The amount of delay between the present sample and the delayed sample is called the differential delay and is denoted as M by most authors. **Figure 10.39** shows the operation of the comb.

An actual CIC filter is composed of multiple integrator-comb sections that are cascaded. A CIC filter has exactly the same number of integrators as combs. According to the associative property you can rearrange the order of the additions in a sequence and the result of the sequence does not change (a + (–b) + c + d + (–e) + (–f) is identical to a + c + d – b – e – f). A CIC filter with a sample rate change uses that property to group all of the integrators together and to group all of the combs together. We place either a downsample or upsample rate changer between the combs and integrators. (See the **DSP and SDR Fundamentals** chapter for a discussion of sample rate changing.)

Figure 10.40 shows that a decimator is an integrator section followed by a down rate change, which is then followed by a comb section. An interpolator turns the system around and puts the comb section first, followed by an up rate changer, which is followed by an integrator section. It is very useful for a hardware implementation that the number of integrators and combs is independent (within reason) from the rate change and that, in general, you can rearrange the inputs, outputs, and rate change to create a decimator and interpolator with the same blocks.

There are three parameters that affect the implementation of a CIC filter. "R" is the up sample rate or down sample rate. "M" is the delay in the comb section and is almost always either one or two. "N" is the number of stages in the comb section (which is required to be the same as the number of integrator sections). The simplification of separating the combs into one section and the integrators into a second section is advantageous for the speed required of the storage elements for the combs. The combs always operate at the low frequency end of the system and the integrators work at the high speed side of the system.

Notice that the associative property would also allow a decimator with the comb first and the integrator after the down converter. Either configuration will give the same results. The reason we always put the comb on the low sample rate side of the system is pragmatic. The comb requires one additional storage element (for the usual M = 1 situation) over what is required for an integrator. Each storage element consumes power when it is clocked. A faster clock and more storage registers translate directly into additional power dissipation. The additional power is an issue when implementing a CIC filter in hardware such as an FPGA.

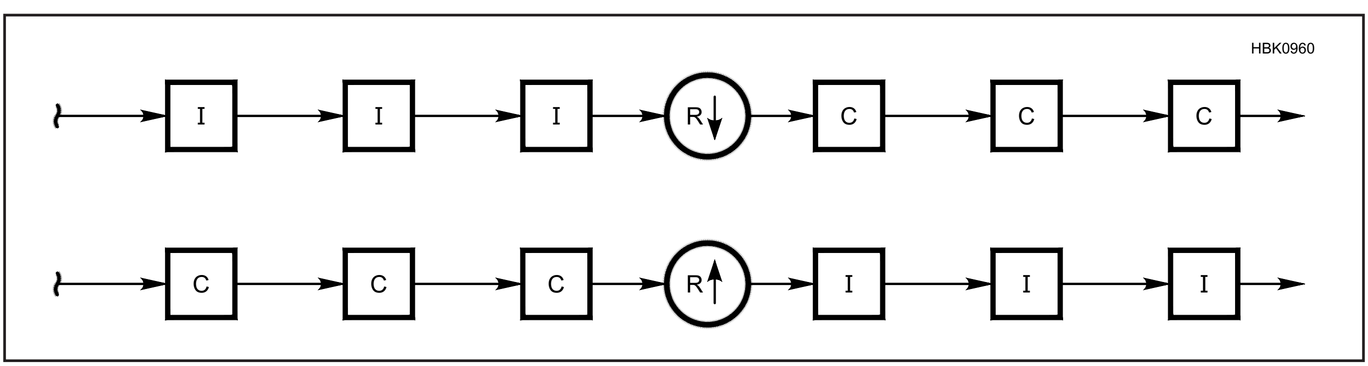

Figure 10.40 — The top CIC filter shows a decimator, where the sample rate is reduced by a value "R." The bottom CIC filter shows an interpolator where the sample rate is increased by a value "R." Note that the difference between the two is the order of the combs and integrators as well as the direction of the rate change. The combs are on the low sample rate side of both systems.

10.5.4 Adaptive Filters

An *adaptive filter* is one that automatically adjusts its filter coefficients under the control of some algorithm. This is often done in situations where the filter characteristics are not known in advance. For example, an adaptive channel equalizer corrects for the non-flatness in the amplitude and phase spectrum of a communications channel due to multipath propagation. Typically, the transmitting station periodically sends a known sequence of data, known as a *training sequence*, which is used by the receiver to determine the channel characteristics and adjust its filter coefficients accordingly.

Another example is an automatic notch filter. An algorithm determines the frequency of an interfering tone and automatically adjusts the notch frequency to remove the tone. Noise cancellation is another application. It can be thought of as the opposite of a notch filter. In this case, all the sine-wave tones in the input signal are considered to be desired and the filter coefficients are configured to enhance them. That method works not only for CW signals but for voice as well since the human voice consists largely of discrete frequencies.

A generic block diagram of an adaptive filter is shown in **Figure 10.41**. The variable filter is typically an FIR type with coefficients calculated by the update algorithm. By some means, an estimate of the desired, unimpaired signal, d, is generated and compared to the filter output y. The difference between y and d is the error, e, which is used by the update algorithm to modify the filter coefficients to improve the accuracy of y. The algorithm is capable of acting as a noise-reduction filter and a notch filter simultaneously. Assuming d is in the form of a pure tone (sine wave), the tone is simultaneously optimized in the y output and minimized in the e output.

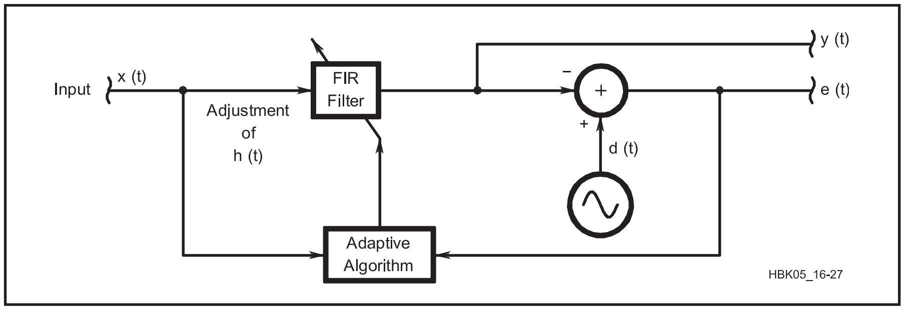

Figure 10.41 — An adaptive filter configuration.

A common algorithm for minimizing the error signal is called *least mean squares* (LMS). The LMS algorithm includes a performance parameter, μ, which can be adjusted between 0 and 1 to control the tradeoff between adjustment speed and accuracy. A value near 1 results in fast convergence but the convergence is not very accurate. For better accuracy at the cost of slower adjustment, lower the value of μ. Some implementations adjust μ on the fly, using a large value at first to get faster lock-in when the error is large then a smaller value after convergence to reduce the error. That works as long as the signal characteristics are not changing too rapidly.

10.6 Quartz Crystal Filters

(The sections on designing and building crystal filters from previous editions are available as a PDF document in the this book's online material.)

Inductor Q values effectively limit the minimum bandwidth that can be achieved with LC band-pass filters. Higher-Q circuit elements, such as quartz crystal, PLZT ceramic and constant-modulus metal alloy resonators, are required to extend these limits. Quartz crystals offer the highest Q and best stability with time and temperature of all available resonators. They are manufactured for a wide range of frequencies from audio to VHF, using cuts (crystal orientations) that suit the frequency and application of the resonator. The AT cut is favored for HF fundamental and VHF overtone use, whereas other cuts (DT, SL and E) are more convenient for use at lower frequencies.

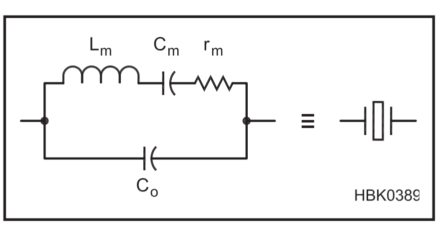

Figure 10.42 — Equivalent circuit of a quartz crystal and its circuit symbol.

Each crystal plate has several modes of mechanical vibration. These can be excited electrically thru the piezoelectric effect, but generally resonators are designed so as to maximize their response on a particular operating frequency using a crystal cut that provides low loss and a favorable temperature coefficient. Consequently, for filter design, quartz crystal resonators are modeled using the simplified equivalent circuit shown in **Figure 10.42**. Here L_m, C_m and r_m represent the *motional* parameters of the resonator at the main operating frequency — r_m being the loss resistance, which is also known as the *equivalent series resistance*, or ESR. C_o is a combination of the static capacitance formed between the two metal electrodes with the quartz as dielectric ($e_r = 4.54$ for AT-cut crystals) and some additional capacitance introduced by the metal case, base and mount. There is a physical relationship between C_m and the static capacitance formed by the resonator electrodes, but, unfortunately, the added holder capacitance masks this direct relationship causing C_o/C_m to vary from 200 to over 500. However, for modern fundamental AT-cut crystals between 1 and 30 MHz, their values of C_m are typically between 0.003 and 0.03 pF. Theoretically, the motional inductance of a quartz crystal should the same whether it is operated on the fundamental or one of its overtones, making the motional capacitance at the nth overtone $1/n^2$ of the value at the fundamental. However, this relationship is modified by the effect of the metal electrodes deposited on either side of the crystal plate and in practice the motional inductance increases with the frequency of the overtone. This makes the motional capacitance at the overtone substantially less than C_m/n^2.

An important parameter for crystal filter design is the unloaded Q at f_s, the resonant frequency of the series arm. This is usually denoted by Q_U.

$$Q_u = 2 \pi f_s L_m/r_m$$

Q_U is very high, often exceeding 100,000 in the lower HF region. Even VHF overtone crystals can have Q_U values over 20,000, making it possible to design quartz crystal band-pass filters with a tremendously wide range of bandwidths and center frequencies.

The basic filtering action of a crystal can be seen from **Figure 10.43**, which shows a plot of attenuation vs frequency for the test circuit shown in the inset**.** The series arm of the crystal equivalent circuit forms a series-tuned circuit, which passes signals with little attenuation at its resonant frequency, f_s, but appears inductive above this frequency and

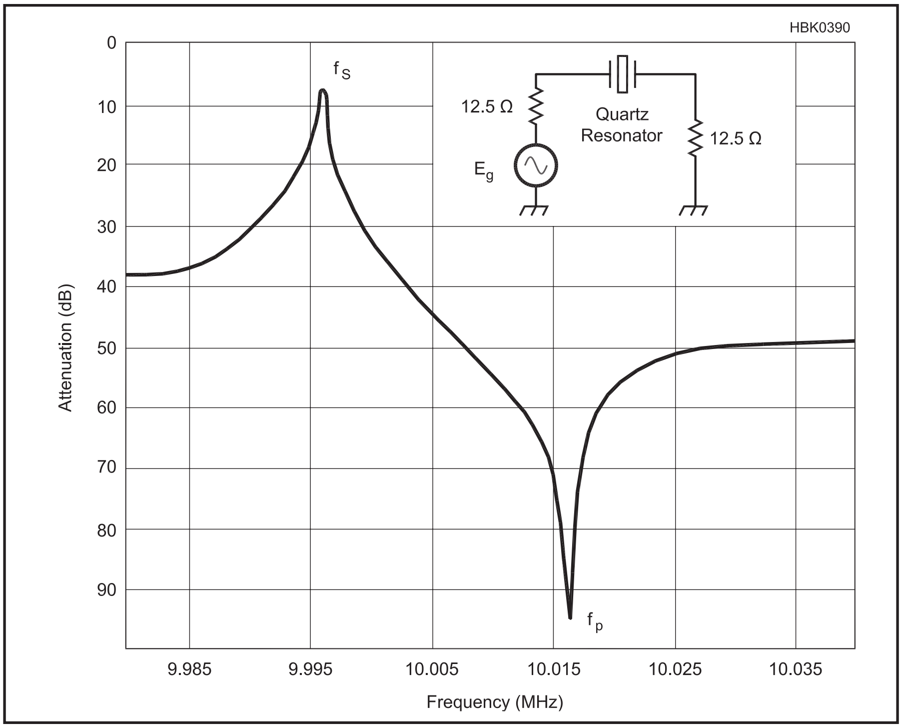

Figure 10.43 — Response of 10 MHz crystal (Cm = 0.0134 pF, Lm = 18.92 mH, ESR = 34 Ω) in a series test circuit (see inset) showing peak of transmission (lowest attenuation) at the series resonant frequency, f_s, and a null (maximum attenuation) caused by the parallel resonance due to C_o at f_p.

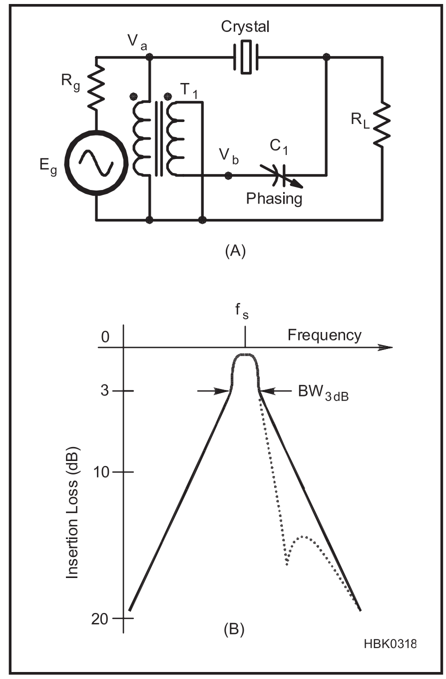

Figure 10.44 — Classic single-crystal filter in A has the response shown in B. The phasing capacitor can be adjusted to balance out C_o (solid line), or set to a lesser or greater value to create a movable null to one side, or other, of the passband (dotted line).

parallel resonates with C_o at f_p to produce a deep notch in transmission. The difference between f_s and f_p is known as the pole-zero separation, or PZ spacing, and is dependent on C_m/C_o as well as f_s. Further information on quartz crystal theory and operation can be found in the **Oscillators and Synthesizers** chapter and in the reference for Bottom.

A simple crystal filter developed in the 1930s is shown in **Figure 10.44**. The voltage-reversing transformer T1 was usually an IF transformer, but nowadays could be a bifilar winding on a ferrite core. Voltages V_a and V_b have equal magnitude but 180° phase difference. When $C_1 = C_o$, the effect of C_o will disappear and a well-behaved single resonance will occur as indicated by the solid line in Figure 10.44B. However, if C_1 is adjusted to unbalance the circuit, a transmission zero (notch) is produced well away from the pass band and by increasing the amount of imbalance this can be brought back toward the edge of the pass band to attenuate close-by interfering CW signals. If C_1 is reduced in value from the balanced setting, the notch comes in from the high side and if C_1 is increased, it comes in from the low side. The dotted curve in Figure 10.44B illustrates how the notch can be set with C_1 less than C_o to suppress adjacent signals just above the pass band. In practice a notch depth of up to 60 dB can be achieved. This form of "crystal gate" filter, operating at 455 kHz, was present in many high-quality amateur communications receivers from the 1930s through the 1960s. When the filter was switched into the receiver IF amplifier the bandwidth was reduced to a few hundred Hz for CW reception. The close-in range of the notch was sometimes improved by making C_1 part of a differential capacitor that could add extra capacitance to either the C_1 or C_o side of the IF transformer. This design could also be used to good effect at frequencies up to 1.7 MHz with an increased minimum bandwidth. However, any crystal gate requires considerable additional IF filtering to achieve a reasonable ultimate attenuation figure, so it should not be the only form of selectivity used in an IF amplifier.

The half-lattice filter shown in **Figure 10.45** offers an improvement in performance over a single-crystal filter. The quartz crystal static capacitors, C_o, cancel each other. The remaining series-resonant arms, if offset in frequency and terminated properly, will produce an approximate 2-pole Butterworth or Chebyshev response. The crystal spacing for simple Chebyshev designs is usually around two-thirds of the bandwidth. Half-lattice filter sections can be cascaded to produce composite filters with multiple poles. Many of the older commercial filters are coupled half-lattice types using 4, 6 or 8 crystals, and this

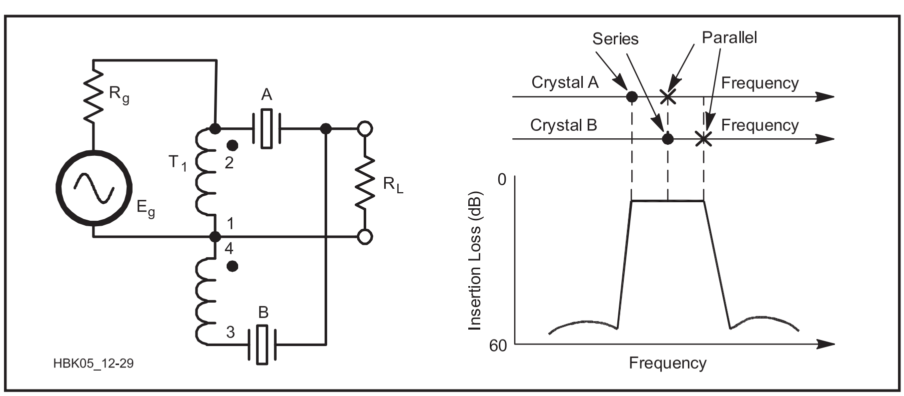

Figure 10.45 — A half-lattice crystal filter. The C_o of one crystal can be made to balance the C_o of the other, or C_o across the higher crystal can be deliberately increased to create nulls on either side of the passband.

Figure 10.46 — A dual monolithic crystal filter has two sets of electrodes acoustically coupled to provide a two-pole response in a single crystal unit. The center lead is connected to the case and grounded in normal operation.

is still the favored technology for some crystal filters at lower frequencies. Very often extra capacitance was added across one crystal in each half lattice to unbalance C_o and provide deep transmission zeroes on either side of the pass band to sharpen up the close-in response at the expense of the attenuation further out. The reference for Steder and Hardcastle discusses the computer design of half-lattice filters.

Most commercial HF and VHF crystal filters produced today use dual monolithic structures as their resonant elements. These are a single quartz plate onto which two sets of metal electrodes have been deposited, physically separated to control the acoustic (mechanical) coupling between them. An example of a dual monolithic filter (2-pole) is shown in **Figure 10.46.** These are available with center frequencies from 9 MHz to well over 120 MHz. In effect, the dual monolithic structure behaves just like a pair of coupled crystal resonators. There is a subtle difference, however, because the static capacitance in duals appears across the input and output terminations and doesn't produce a null above the pass band, as it would for two electrically coupled crystal resonators. Multi-pole filters can be built by coupling duals together using external capacitors, the C_o from each dual input and output being absorbed in the coupling capacitors or terminations, or by including more acoustically coupled resonators in the electrode structure on a single quartz plate. Though not common, 3-pole and 4-pole monolithic filters housed in standard, single crystal holders are available from some manufacturers.

10.6.1 Filter Parameters

An ideal band-pass filter would pass the desired signal with no loss and completely attenuate everything else. Practically, it is not possible to achieve such a response with a finite number of resonators, and approximations to this ideal have to be accepted. Greater stop-band attenuation and steeper sides can be achieved if more and more crystals are used, and the response gets nearer to the ideal "brickwall" one. This feature of a filter is expressed as a *shape factor*, which specifies the ratio of the bandwidth at an attenuation of 60 dB to the bandwidth at 6 dB — both these levels being taken relative to the actual pass band peak to eliminate insertion loss from the calculation. An ideal brickwall filter would have a shape factor of 1, and practical filters have shape factors that depend on the number of crystals used in the design and the type of response chosen for the pass band. A 1 dB Chebyshev design, for example, can typically produce shape factors that vary from about 4.1 for a 4-pole to 1.5 for a 10-pole, but the actual figures obtained in practice are very much dependent on Q_u and how much greater it is than the filter Q, defined by f_o/BW, where f_o is the center frequency and BW the 3 dB bandwidth. The ratio of Q_u to f_o/BW is often quoted as Q_o, or q_o, and this, along with the order and type of response, determines the insertion loss of a crystal filter. Q_o also determines how closely the pass band follows the design response, and how much the passband ripple is smoothed out and the edges of the response rounded off by crystal loss.

Commercial filter manufacturers usually choose the Chebyshev equiripple design for SSB, AM and FM bandwidths because it gives the best compromise between passband response and the steepness of the sides, and 1 dB-ripple Chebyshev designs are pretty standard for speech bandwidths. Tolerances in component values and crystal frequencies can cause the ripple in the pass band to exceed 1 dB, so often the maximum ripple is specified as 2 dB even though the target ripple is lower. Insertion loss, the signal loss going thru the filter, also varies with the type of response and increases as the order of the filter increases. The insertion loss for a given order and bandwidth is higher for high-ripple Chebyshev designs than it is for low-ripple ones, and Butterworth designs have lower insertion loss than any Chebyshev type for a given bandwidth. Pass-band amplitude response and shape factor are important parameters for assessing the performance of filters used for speech communications.

However, group delay is also important for data and narrow bandwidth CW reception. Differential group delay can cause signal distortion on data signals if the variations are greater than the automatic equalizer can handle. Ringing can be an annoying problem when using very narrow CW filters, and the group delay differential across the pass band must be minimized to reduce this effect. When narrow bandwidth filters are being considered, shape factor has to be sacrificed to reduce differential group delay and its associated ringing problems. It's no good having a narrow Chebyshev design with a shape factor of 2 if the filter produces unacceptable ringing and is intolerable to use in practice.

Both Bessel and linear phase (equiripple 0.05°) responses have practically constant, low group delay across the entire pass band and well beyond on either side, making either a good choice for narrow CW or specialized data use. They also have the great advantage of offering the lowest possible insertion loss of all the types of response currently in use, which is important when Q_o is low, as it often is for very narrow bandwidth filters. The insertion loss of the Bessel design is marginally lower than that of the linear phase, but the latter has a superior shape factor giving it the best balance of low group delay and good selectivity. A 6-pole linear phase (equiripple 0.05°) design has a shape factor of 3.39, whereas a 6-pole Bessel has 3.96. The Gaussian-to-12 dB response has a better shape factor than that of either the linear phase or Bessel designs but the group delay across the pass band is not as flat and pronounced peaks (ears) are beginning to appear at the band edges with a 6-pole design. The Gaussian-to-6 dB group delay is reasonably flat for 3- and 4-pole filters, but significant ears appear toward the passband edges in designs with 5 poles or more.

10.6.2 Crystal Filter Evaluation

The simplest means of assessing the performance of a crystal filter is to temporarily install it in a finished transceiver, or receiver, and use a strong on-air signal, or locally generated carrier, to run thru the filter pass band and down either side to see if there are any anomalies. Provided that the filter crystals have been carefully characterized in the first instance, and computer modeling has shown that the design is close to what's required, this may be all that's required to confirm a successful project. However, more elaborate checks on both the pass band and stop band can be made if a DDS signal generator, or vector network analyzer (VNA), is available. A test set-up for evaluating the response of a filter using a DDS generator requires an oscilloscope to display the response, whereas a VNA controlled via the USB port of a PC can display the response on the PC screen with suitable software — see the **Test Equipment and Measurements** chapter. In addition, it can measure the phase and work out the group delay of the filter. VNAs can also be used to characterize the crystals prior to making the filter as well as evaluating filter performance after completion.

10.7 SAW Filters

The resonators in a monolithic crystal filter are coupled together by bulk acoustic waves. These acoustic waves are generated and propagated in the interior of a quartz plate. It is also possible to launch, by an appropriate transducer, acoustic waves that propagate only along the surface of the quartz plate. These are called "surface-acoustic-waves" because they do not appreciably penetrate the interior of the plate.

A *surface-acoustic-wave* (SAW) filter consists of thin aluminum electrodes, or fingers, deposited on the surface of a piezoelectric substrate as shown in **Figure 10.47**. Lithium niobate ($LiNbO_3$) is usually favored over quartz because it yields less insertion loss. The electrodes make up the filter's transducers. RF voltage is applied to the input transducer and generates electric fields between the fingers. The piezoelectric material vibrates in response, launching an acoustic wave along the surface. When the wave reaches the output transducer it produces an electric field between the fingers. This field generates a voltage across the load resistor.

Since both input and output transducers are not entirely unidirectional, some acoustic power is lost in the acoustic absorbers located behind each transducer. This lost acoustic power produces a mid-band electrical insertion loss typically greater than 10 dB. The SAW filter frequency response is determined by the choice of substrate material and finger pattern. The finger spacing, (usually one-quarter wavelength) determines the filter center frequency. Center frequencies are available from 20 to 1000 MHz. The number and length of fingers determines the filter loaded Q and shape factor.

Loaded Qs are available from 2 to 100, with a shape factor of 1.5 (equivalent to a dozen poles). Thus the SAW filter can be made broadband much like the LC filters that it replaces. The advantage is substantially reduced volume and possibly lower cost. SAW filter research was driven by military needs for exotic amplitude-response and time-delay requirements. Low-cost SAW filters are presently found in television IF amplifiers where high mid-band loss can be tolerated.

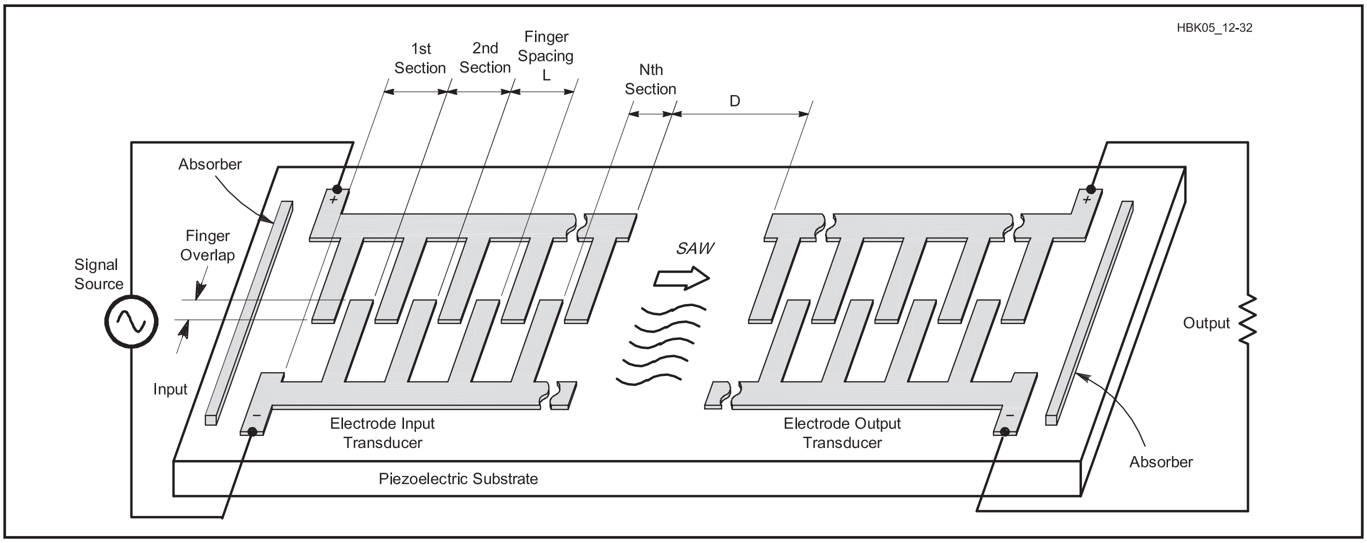

Figure 10.47 — The *interdigitated* transducer, on the left, launches SAW energy to a similar transducer on the right (see text).

10.8 Transmission Line VHF/UHF/Microwave Filters

LC filter calculations are based on the assumption that the reactances are *lumped*— that the physical dimensions of the components are considerably less than the operating wavelength. In such cases the unavoidable inter-turn parasitic capacitance associated with inductors and the unavoidable series parasitic inductance associated with capacitors are neglected as being secondary effects. If careful attention is paid to circuit layout and miniature components are used, lumped LC filter technology can be used up to perhaps 1 GHz.

Replacing lumped reactances with selected short sections of *Transverse Electromagnetic Mode* (TEM) transmission lines results in transmission line filters. (In TEM the electric and magnetic fields associated with a transmission line are at right angles (transverse) to the direction of wave propagation.) Coaxial cable, stripline and microstrip are examples of TEM components. Waveguides and waveguide resonators are not TEM components.

Coaxial cable transmission line filters are often used at HF and VHF frequencies. Stripline and microstrip transmission-line filters predominate from 500 MHz to 10 GHz. In addition they are often used down to 50 MHz when narrowband ($Q_L > 10$) band-pass filtering is required. In this application they exhibit considerably lower loss than their LC counterparts and are useful at frequencies where coaxial transmission lines are too lossy. A detailed treatment of the use of coaxial cable to form transmission line filters is presented in the **Transmission Lines** chapter. This section focuses on stripline and microstrip filters used at VHF and above.

10.8.1 Stripline and Microstrip Filters

Figure 10.48 shows three popular transmission lines used in transmission line filters. The circular coaxial transmission line (*coax*) shown in Figure 10.48A consists of two concentric metal cylinders separated by dielectric (insulating) material. The first transmission-line filters were built from sections of coaxial line. Their mechanical fabrication is expensive and it is difficult to provide electrical coupling between line sections.

Fabrication difficulties are reduced by the use of shielded strip transmission line (*stripline*) shown in Figure 10.48B. The outer conductor of stripline consists of two flat parallel metal plates (ground planes) and the inner conductor is a thin metal strip. Sometimes the inner conductor is a round metal rod. The dielectric between ground planes and strip can be air or a low-loss plastic such as polyethylene. The outer conductors (ground planes or shields) are separated from each other by distance *b*.

Striplines can be easily coupled together by locating the strips near each other as shown in Figure 10.48B. Stripline Z_0 vs width (*w*) is plotted in **Figure 10.49**. Air-dielectric stripline technology is best for low bandwidth ($Q_L > 20$) band-pass filters.

The most popular transmission line at UHF and microwave is *microstrip* (unshielded stripline), shown in Figure 10.48C. It can be fabricated with standard printed-circuit processes and is the least expensive configuration. In microstrip the outer conductor is a single flat metal ground-plane. The inner conductor is a thin metal strip separated from the ground-plane by a solid dielectric substrate. Typical substrates are 0.062 inch G-10 fiberglass ($\varepsilon = 4.5$) for the 50 MHz to 1 GHz frequency range and 0.031 inch Teflon ($\varepsilon = 2.3$) for frequencies above 1 GHz. Unfortunately, microstrip has the most loss of the three types of transmission line; therefore it is not suitable for narrow, high-Q, band-pass filters.

Conductor separation must be minimized or radiation from the line and unwanted coupling to adjacent circuits may become problems. Microstrip characteristic impedance and the effective dielectric constant (ε) are shown in **Figure 10.50**. Unlike coax and stripline, the effective dielectric constant is less than that of the substrate since a portion of the electromagnetic wave propagating along the microstrip travels in the air above the substrate.

The characteristic impedance for stripline and microstrip-lines that results in the lowest loss is not 75 Ω as it is for coax. Loss decreases as line width increases, which leads to clumsy, large structures. Therefore, to conserve space, filter sections are often constructed from 50 Ω stripline or microstrip stubs even though the loss at that characteristic impedance is not a minimum for that type of transmission line.

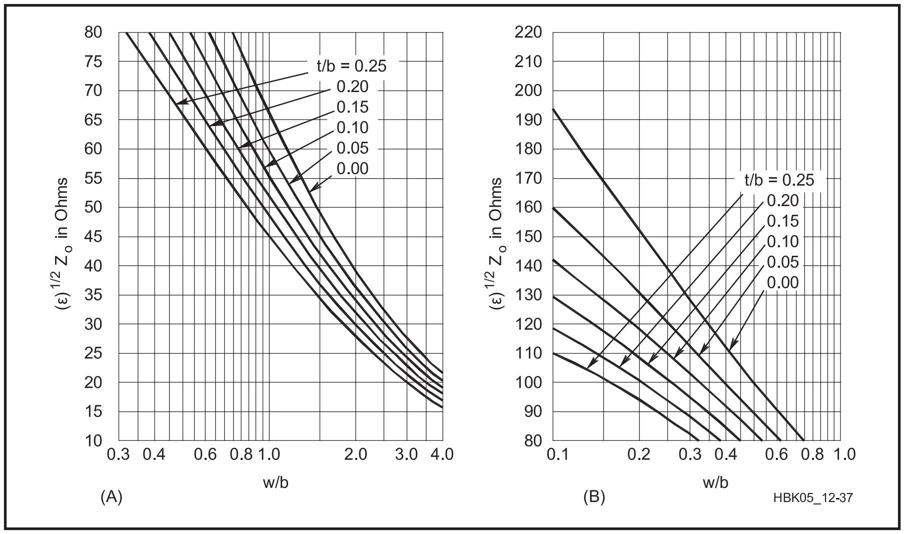

Figure 10.49 — The Z_0 of stripline varies with *w*, *b* and *t* (conductor thickness). See Figure 10.48B. The conductor thickness is *t* and the plots are normalized in terms of *t/b*.

10.8.2 Transmission Line Band-Pass Filters

Band-pass filters can also be constructed from transmission line stubs. (See the **Transmission Lines** chapter for information on stub behavior and their use as filters at HF and VHF.) At VHF the stubs can be considerably shorter than a quarter-wavelength (¼ λ), yielding a compact filter structure with less mid-band loss than its LC counterpart. The single-stage 146 MHz stripline band-pass filter shown in **Figure 10.51** is an example.

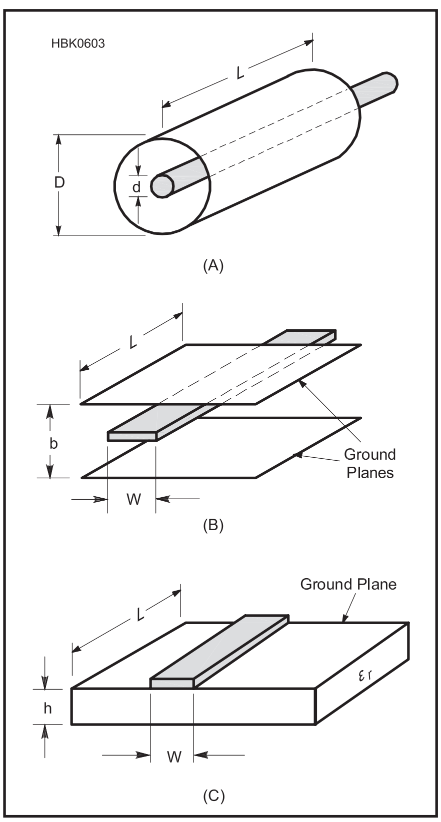

Figure 10.48 — Transmission lines. A: Coaxial line. B: Coupled stripline, which has two ground planes. C: Microstrip has only one ground plane.

Z_0 Ω	ε =1 (AIR) W/h	ε =2.3 (RT/Duroid) W/h	$\sqrt{\varepsilon_e}$	ε =4.5 (G−10) W/h	$\sqrt{\varepsilon_e}$
25	12.5	7.6	1.4	4.9	2.0
50	5.0	3.1	1.36	1.8	1.85
75	2.7	1.6	1.35	0.78	1.8
100	1.7	0.84	1.35	0.39	1.75
	$\sqrt{\varepsilon}$ =1				

HBK05_12-38

Figure 10.50 — Microstrip parameters (after H. Wheeler, *IEEE Transactions on MTT*, March 1965, p 132). e_e is the effective ε.

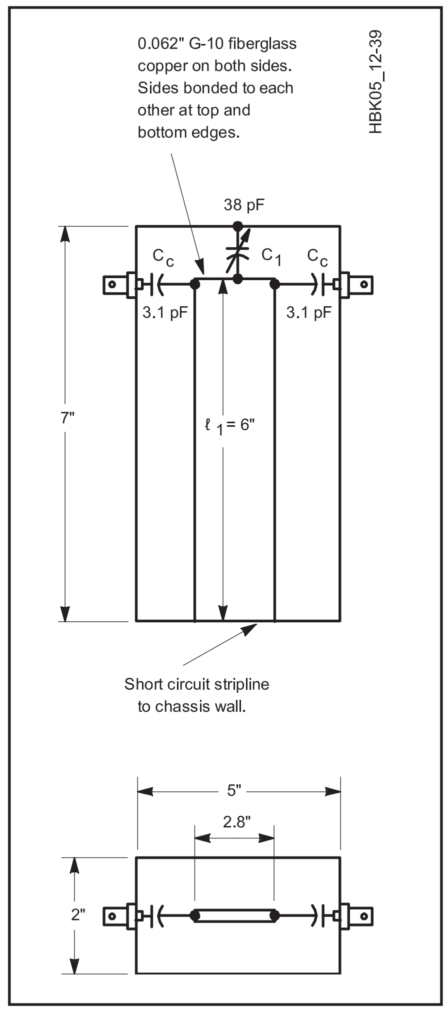

Figure 10.51 — This 146 MHz stripline band-pass filter has been measured to have a Q_L of 63 and a loss of approximately 1 dB.

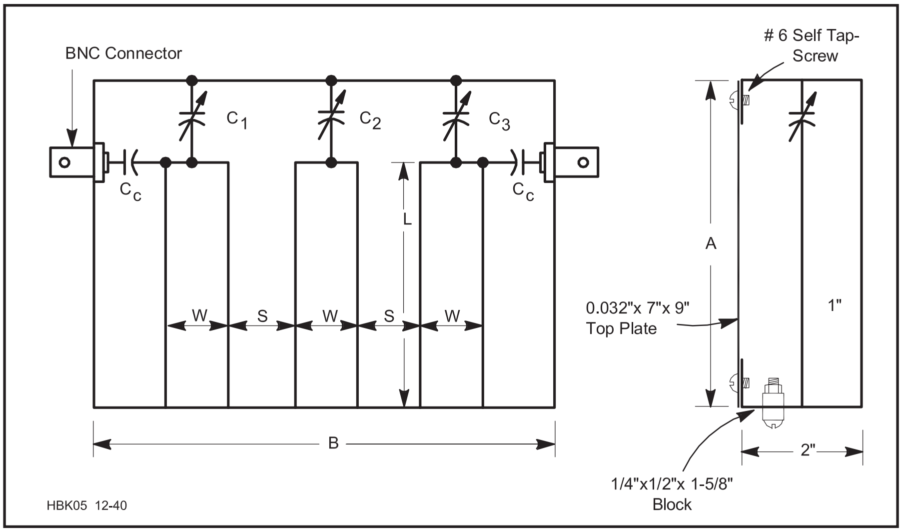

Figure 10.52 — This Butterworth filter is constructed in combline. It was originally discussed by R. Fisher in December 1968 *QST*.

This filter consists of a single inductive 50 Ω strip-line stub mounted into a 2 × 5 × 7 inch aluminum box. The stub is resonated at 146 MHz with the "APC" variable capacitor, C1. Coupling to the 50 Ω generator and load is provided by the coupling capacitors C_c. The measured performance of this filter is: f_o = 146 MHz, BW = 2.3 MHz (Q_L = 63) and mid-band loss = 1 dB.

Single-stage stripline filters can be coupled together to yield multistage filters. One method uses the capacitor coupled band-pass filter synthesis technique to design a 3-pole filter. Another method allows closely spaced inductive stubs to magnetically couple to each other. When the coupled stubs are grounded on the same side of the filter housing, the structure is called a "combline filter." Three examples of combline band-pass filters are shown in **Figure 10.52** and a set of VHF/UHF filter designs by W1GHZ is included in the projects section of this chapter. These filters are constructed in 2 × 7 × 9 inch aluminum boxes. The article describing these filters by Reed Fisher, W2CQH, is available in this book's online information and is listed in the references, as well.

10.8.3 Quarter-Wave Transmission Line Filters

The reactance of a ¼ λ shorted-stub is infinite, as discussed in the **Transmission Lines** chapter. Thus, a ¼ λ shorted stub behaves like a parallel-resonant LC circuit. Proper input and output coupling to a ¼ λ resonator yields a practical band-pass filter. Closely spaced ¼ λ resonators will couple together to form a multistage band-pass filter. When the resonators are grounded on opposite walls of the filter housing, the structure is called an *interdigital filter* because the resonators look like interlaced fingers. Two examples of 3-pole UHF interdigital filters are shown in **Figure 10.53**. Design graphs for round-rod interdigital filters are given in the reference for Metcalf. The ¼ λ resonators may be tuned by physically changing their lengths or by tuning the screw opposite each rod.

If the short-circuited ends of two ¼ λ resonators are connected to each other, the resulting ½ λ stub will remain in resonance, even when the connection to the ground-plane is removed. Such a floating ½ λ microstrip line, when bent into a U-shape, is called a *hairpin* resonator. Closely coupled hairpin resonators can be arranged to form multistage band-pass filters. Microstrip hairpin band-pass filters are popular above 1 GHz because they can be easily fabricated using photo-etching techniques. No connection to the ground-plane is required.

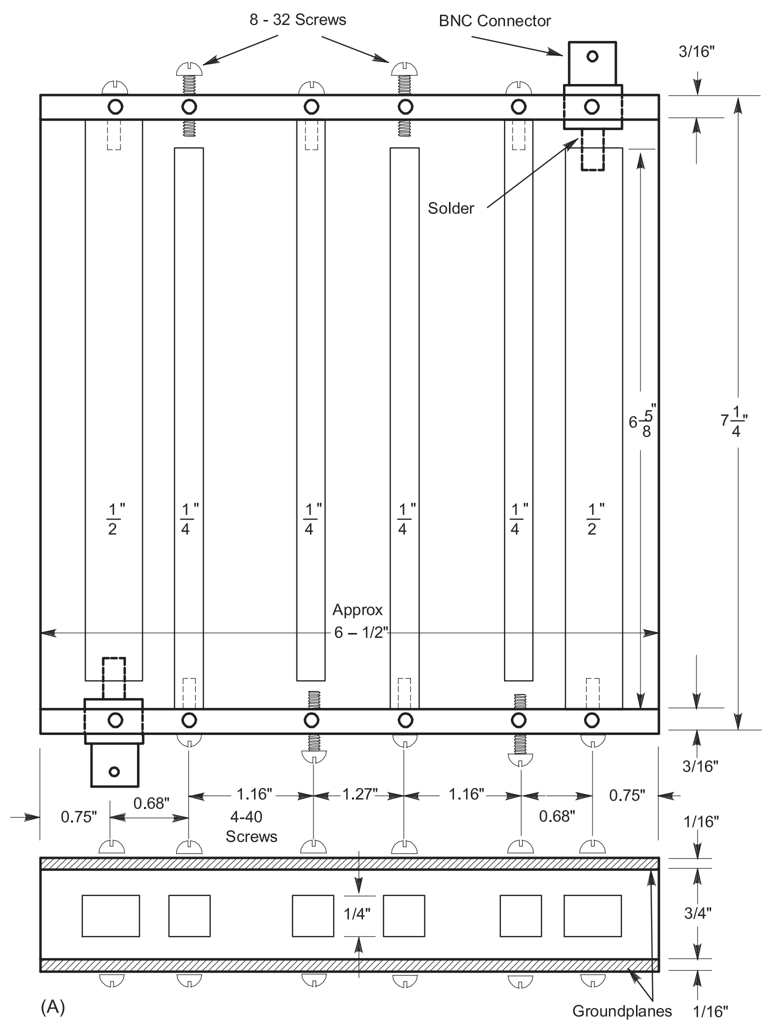

Figure 10.53 — These 3-pole Butterworth filters for 432 MHz (shown at A, 8.6 MHz bandwidth, 1.4 dB passband loss) and 1296 MHz (shown at B, 110 MHz bandwidth, 0.4 dB passband loss) are constructed as interdigitated filters. The material is from R. E. Fisher, March 1968 *QST*.

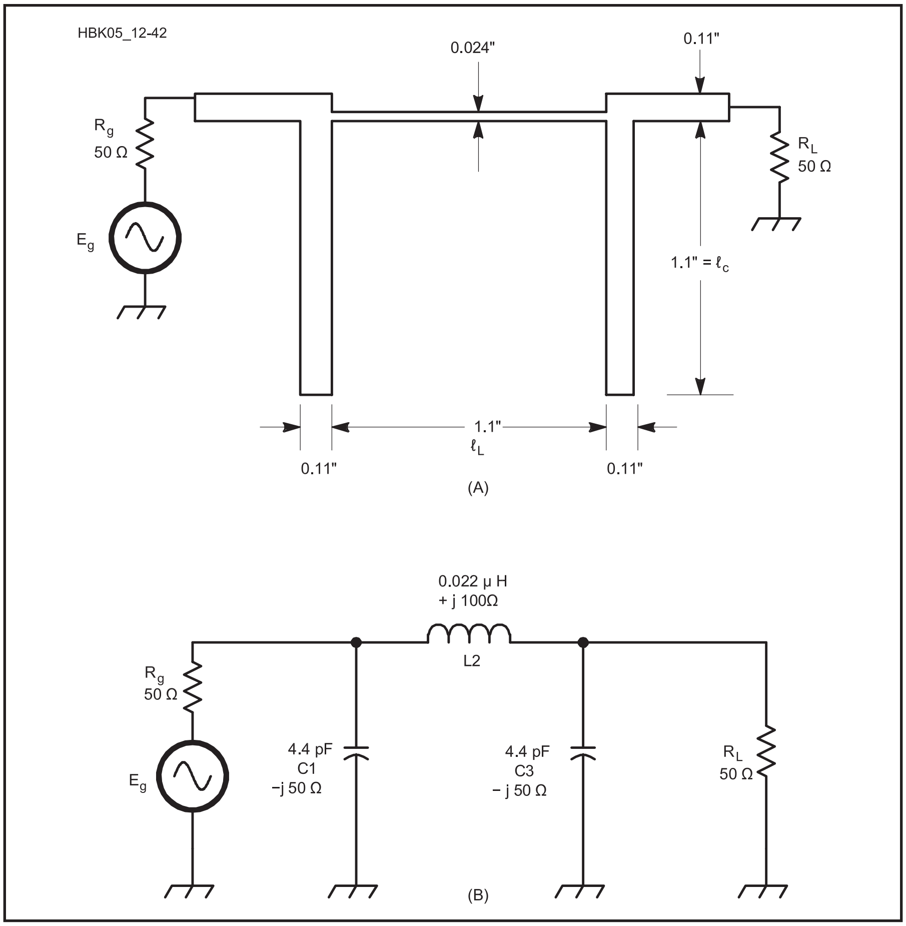

Figure 10.54 — A microstrip 3-pole emulated-Butterworth low-pass filter with a cutoff frequency of 720 MHz. A: Microstrip version built with G-10 fiberglass board (ε = 4.5, h = 0.062 inches). B: Lumped LC version of the same filter. To construct this filter with lumped elements very small values of L and C must be used and stray capacitance and inductance would have to be reduced to a tiny fraction of the component values.

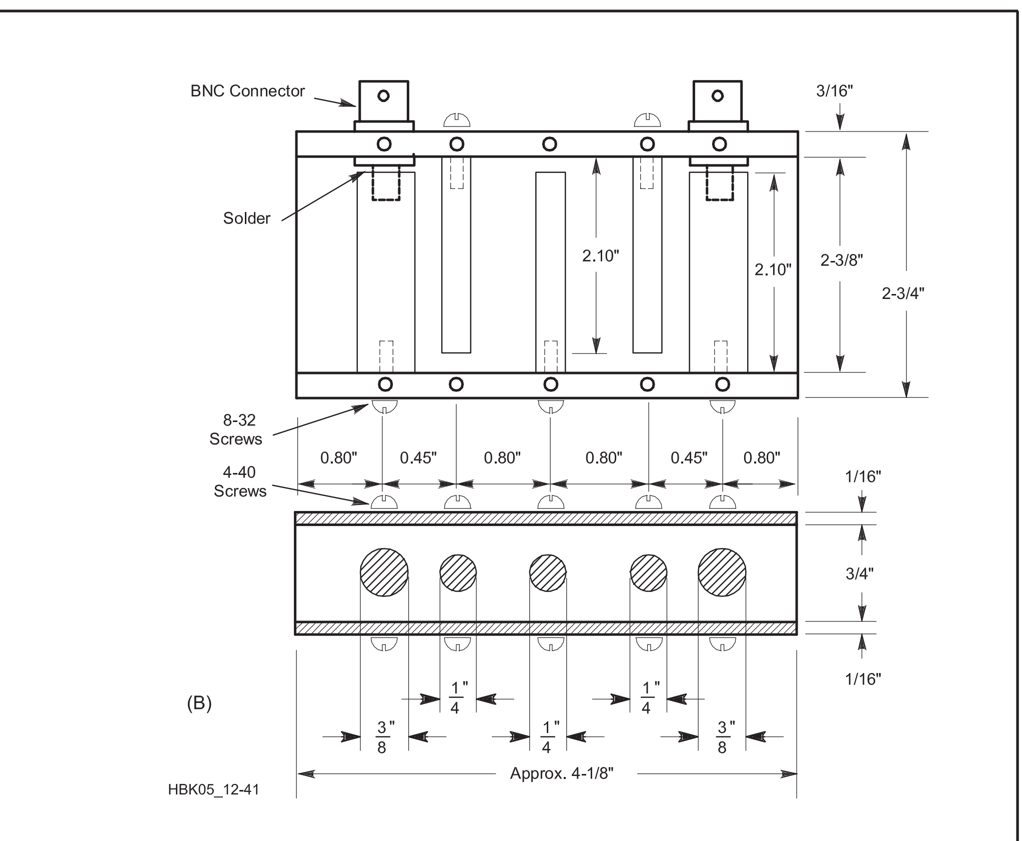

10.8.4 Emulating LC Filters with Transmission Line Filters

Low-pass and high-pass transmission-line filters are usually built from short sections of transmission lines (stubs) that emulate lumped LC reactances. Sometimes low-loss lumped capacitors are mixed with transmission line inductors to form a hybrid component filter. For example, consider the 720 MHz, 3-pole microstrip low-pass filter shown in **Figure 10.54A** that emulates the LC filter shown in Figure 10.54B. C1 and C3 are replaced with 50 Ω open-circuit shunt stubs ℓ_C long. L2 is replaced with a short section of 100-Ω line ℓ_L long. The LC filter, Figure 10.54B, was designed for f_c = 720 MHz. Such a filter could be connected between a 432 MHz transmitter and antenna to reduce harmonic and spurious emissions. A reactance chart shows that X_C is 50 Ω, and the inductor reactance is 100 Ω at f_c. The microstrip version is constructed on G-10 fiberglass 0.062 inch thick, with ε = 4.5. Then, from Figure 10.50, *w* is 0.11 inch and ℓ_C = 0.125 l_g for the 50 Ω capacitive stubs. Also, from Figure 10.50, *w* is 0.024 inch and ℓ_L is 0.125 l_g for the 100-Ω inductive line. The inductive line length is approximate because the far end is not a short circuit. l_g is 300/(720 × 1.75) = 0.238 m, or 9.37 inches Thus ℓ_C is 1.1 inch and ℓ_L is 1.1 inches.

This microstrip filter exhibits about 20 dB of attenuation at 1296 MHz. Its response rises again, however, around 3 GHz. This is because the fixed-length transmission line stubs change in terms of wavelength as the frequency rises. This particular filter was designed to eliminate third-harmonic energy near 1296 MHz from a 432 MHz transmitter and does a better job in this application than the Butterworth filter in Figure 10.53 which has spurious responses in the 1296 MHz band.

10.9 Cavity and Helical Filters

10.9.1 Cavity Filters

The following material is based on the article "Understanding the Cavity Duplexer" by John Portune, W6NBC, available at **w6nbc.com/articles/duplexer.pdf**. A detailed analysis of cavities and other transmission line resonators is contained in *Microwave Engineering* by David Pozar. The website Repeater-Builders (**www.repeater-builder.com**) also contains a lot of additional information about the use of cavities as duplexers for repeater systems. (Also see this book's **Repeater Systems** chapter.)

This section discusses amateur radio's most common type of cavity filter, the ¼-λ coaxial cavity used in repeater duplexers. Overviews of commercial cavity filters used at VHF and UHF from Telewave can be found online by searching for "A Brief Guide to Telewave's Cavity Filters" and "Tips on Navigating Cavity Filter Types, Performance, and Features." A tutorial on cavity construction and tuning is available from TXRX at **www.txrx.com/resources/combiner-basics**.

CAVITY BASICS

The ¼-λ cavity has three important properties: (1) the ability to handle power, (2) high Q, and (3) low loss. It is this unique combination of all three in that has long made the resonant cavity the dominant choice for use as duplexer filters.

A *coaxial cavity resonator* is an open cylinder with highly conductive walls. The cavity's size is such that it resonates at RF. A cavity resonator is analogous to an acoustic organ pipe, a penny whistle, or flute. The mathematical descriptions of electromagnetic and acoustic cavities are nearly identical.

The equivalent circuit of the ¼-λ coaxial cavity is a parallel LC "tank" circuit as shown in **Figure 10.55**. This makes the basic concept of the resonant cavity a little easier to visualize. Inside the outer cylinder is a smaller metal cylinder which forms the coaxial center conductor. See **Figure 10.56**. The center conductor is connected to the outer cylinder at one end of the cavity but not at the other. Figure 10.56 also shows the coupling loops and tuning adjustment which are discussed below.

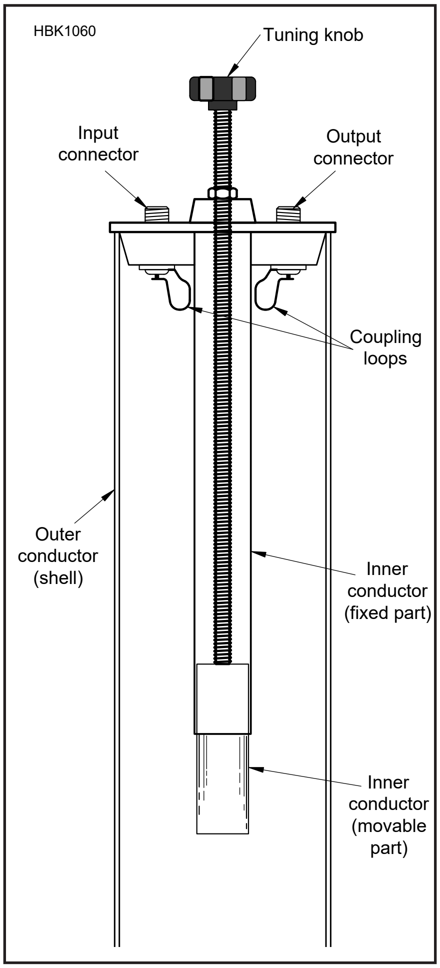

Figure 10.56 — Components of a ¼-λ coaxial cavity. The shorted end is at the top where connectors are located and the open end is effectively at the end of the inner conductor. The bottom of the cavity is covered by a conductive cap.

This configuration is a short length of large-diameter air-insulated coaxial transmission line, shorted at one end. Thus, the common duplexer cavity is a simply a ¼-λ shorted coaxial stub. (See the **Transmission Lines** chapter.) The open end is covered with a conductive cap placed just a small distance away from the open end of the center conductor. The cover has little effect on the action of the filter and keeps the cavity free of dust and other contaminants.

FIELD STRENGTH AND ORIENTATION

The energy in the cavity is in the form of an electromagnetic (EM) field. The electric component of the field (E) is parallel to the length of the cavity, as shown in **Figure 10.57**. The magnetic component of the field (H) is at right angles to the electric component, in concentric circles around the center conductor. The H component and the associated currents in the cavity walls are strongest at the shorted end of the cavity and weaker at the open end. The E component and its associated voltages across the cavity wall and center conductor are strongest at the open end and weakest at the shorted end. Both, however, are stronger near the center conductor and weaker near the outer conductor. The orientation of the field components is important when using loops and probes for coupling signals into and out of the cavity. It is also possible to create *overtone* or *harmonic* modes in which the E and H components take on a more complex structure in the cavity.

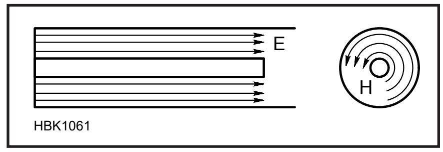

Figure 10.57 — Electric (E) and magnetic (H) field orientation.

CAVITY DIMENSIONS

The outer cylinder's diameter should not be larger than roughly ⅓-λ to avoid the overtone modes which result in higher losses. This means that the approximate diameter limit for 450 MHz cavities is roughly 8 inches and 25 inches for 2 meter cavities.

Optimum Q (for the sharpest filter response) is obtained for a ratio of 3.6 between the outer cylinder's inside diameter and the inner cylinder's outside diameter. This geometry results in the cavity having a characteristic impedance of roughly 77 Ω. The inner surfaces of the cavity are often silver plated for lower loss which also increases Q.

The resonant frequency is determined almost exclusively by the length of the center conductor which is approximately ¼-λ long. For the 70 cm band that is about 6 inches and for 2 meters roughly 19 – 20 inches. On 6 meters and 10 meters, these cavities become very large, making helical resonators (discussed in the following sections) a more practical choice. As shown in Figure 10.56, the center conductor is usually made so that its length is adjustable for tuning. (Note that the sliding contact surface must be carefully designed to avoid resistive losses.)

The length of the outer tube of a ¼-λ cavity has little effect on the tuning, so it is usually made a little longer than the center conductor for mechanical convenience in attaching a cap and in allowing room for adjusting the length of the inner (center) conductor.

10.9.2 Coupling to the Cavity

There are four practical ways to couple to a cavity: loops, probes, ports, and taps. All four

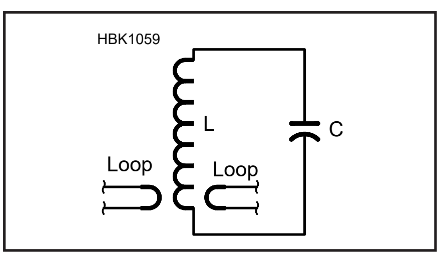

Figure 10.55 — Equivalent circuit of a cavity — parallel LC tank circuit.

coupling methods ultimately perform more or less the same after critical adjustments. However, loops are the most common choice since they are much easier to build and adjust.

LOOP COUPLING

A coupling loop acts as a single-turn coil excited from a connector mounted through the cavity wall. It is most often placed at the shorted end but may also be placed along the side of the cavity. The far end of the loop is connected to the cavity as in Figure 10.56. Loops are analogous to the link windings shown in Figure 10.55.

A loop couples to the H component of the field. Coupling is strongest when the loop is perpendicular to the H component. Since the H component, as we learned, lies in concentric circles around the center conductor, the loop is normally placed parallel to the length of the cavity and along the cavity's radius. It also couples best where the H component is strongest. This is near the shorted end and near the center conductor. Figure 10.56 shows loops in the maximum coupling position.

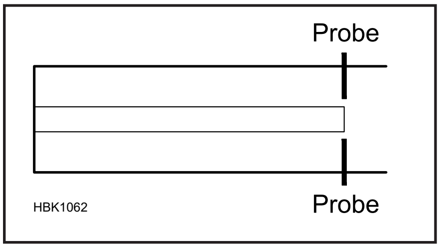

Figure 10.58 — Probe position for maximum coupling.

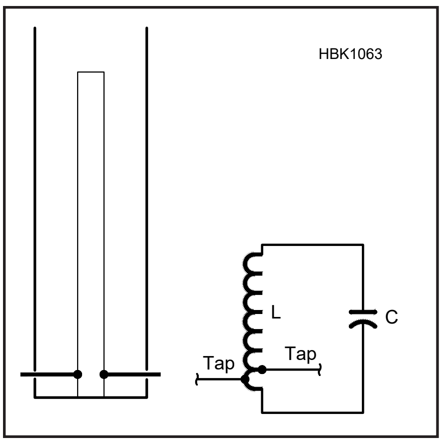

Figure 10.59 — Tap coupling and equivalent circuit.

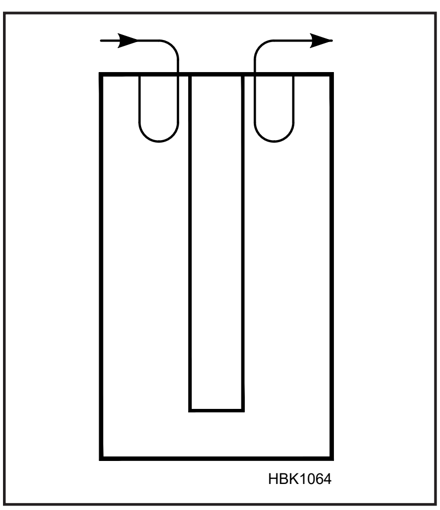

Figure 10.60 — Cavity in series with transmission line. This is a band-pass (BP) configuration.

PROBE COUPLING

A less frequently used coupling method is a probe as shown in **Figure 10.58**. The probe acts as one electrode of a capacitor that couples to the E field. The center conductor of the cavity acts as other half of the coupling capacitor. As opposed to a grounded loop, a probe is open at the end. A probe couples best when it is perpendicular to the E field and is placed where the E field is strongest. This is at the open end of the cavity, near the center conductor.

One disadvantage of probes is that they must be placed at the high voltage (open) end of the cavity. With probes, arcing is a potential problem, even at moderate power levels. In a good quality cavity, Q can easily reach 1000. If 100 W in a 50 Ω feed line produces 71 V, a cavity with a Q of 1000 connected to that feed line will experience 71 × 1000 or 71 kV. Very high voltages can exist at the open end of a cavity where a probe is likely to be installed.

PORT COUPLING

A third way to couple energy, typically between adjacent cavities, is to cut a hole in the outer walls of both cavities to allow some of the field to leak through into the adjacent cavity. This is called *port coupling*. A number of duplexer designs successfully implement this method. Helical resonators, often found in receiver front ends, commonly use port coupling. It is economical and a space saver compared to loops or probes. Small mobile duplexers of rectangular cross section often use port coupling.

The main difficulty with port coupling is mechanical. Varying the position and coupling of loops and probes is easy. To change the amount of port coupling one must physically change the size of the port. This precludes easy experimentation. Since duplexer filters are often made of cylindrical tubing, a port between cavities is also not easy to fabricate and may not be practical for the home builder.

TAP COUPLING

The final method, illustrated in **Figure 10.59**, is *tap coupling*. On the left is an actual cavity with taps and on the right is the equivalent LC circuit with taps. By adjusting the position of the taps, one can achieve a good impedance match as well as efficient coupling. In small cavities, where loops could be too large for the space available, a tap is an easy way to obtain good coupling.

The disadvantage of tap coupling, however, is isolation. If the type of cavity you wish to use requires two ports, an input and an output tap, isolation between ports is difficult to achieve. Tap coupling finds its best application in single-port cavities, such as notch cavities. It is seldom used in band-pass cavities.

10.9.3 Cavity Filter Types

By selecting the coupling method and connection to the transmission line, as ¼-λ cavity can be configured to perform any of three basic filter types: band pass (BP), band-reject or notch (BR), and band-pass/band-reject (BP-BR). A BP cavity *passes* a small band of frequencies and all others are rejected. A BR cavity only rejects a small band of frequencies. All other frequencies are unaffected.

In the series configuration all of the energy

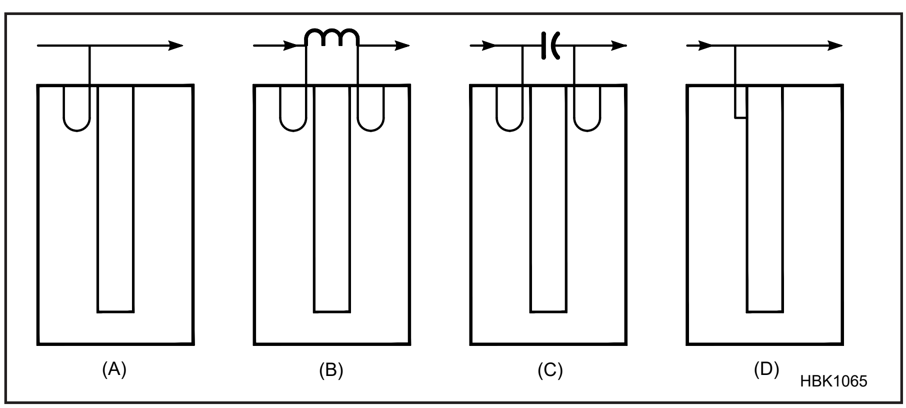

Figure 10.61 — Cavities in parallel with transmission line. These are band-reject or notch (BR) configurations. See text for explanation of different configuration responses.

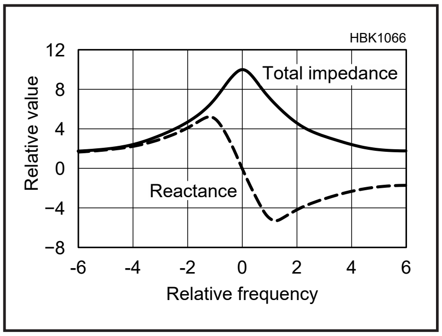

Figure 10.62 — Impedance and reactance of a parallel LC circuit.

passes through the cavity. It is coupled into the cavity by one loop and out by the other as in **Figure 10.60**. (Probes, ports, or taps could also be used.) In a parallel or shunt configuration as in **Figure 10.61**, energy is coupled into and out of the cavity so as to reinforce or cancel the energy in the transmission line.

To understand how each type of filtering is achieved, examine **Figure 10.62** which shows the impedance curve for a parallel LC circuit, the cavity's equivalent circuit. At resonance, impedance reaches a maximum. The absolute value depends on the Q of the cavity. Off-resonance, like a parallel LC circuit, the cavity's total impedance is low.

BAND-PASS (BP) CAVITIES

To obtain a BP response, the cavity must be in series with the transmission line as in Figure 10.60. In this series configuration, all of the energy passes through the cavity. It is coupled into the cavity by one loop and out by the other loop. At the resonant frequency, the high-Q cavity readily accepts the energy supplied to it by the input loop. At the output loop the H field couples energy back into the transmission line. At resonance very little signal is lost. Essentially, to the energy in the transmission line, the cavity is invisible. Away from resonance, at a higher or lower frequency, the cavity couples less energy back into the transmission line. This suppresses off-frequency signals. A typical BP cavity frequency response is shown in **Figure 10.63A**. This cavity has a 2.5 MHz bandwidth at 50 dB below the peak response.

BAND-REJECT (BR) OR NOTCH CAVITIES

By placing the cavity in parallel with the line, as in Figure 10.61, we create a band-reject (BR) or notch cavity or shunt configuration. Parallel or shunt-connected cavities are sometimes called *suck out cavities*. A BR cavity is capable of many dB of signal rejection, far more than a BP cavity. The deep narrow notch shown in Figure 10.63B is why BR cavities are so useful in cavity duplexers.

In Figure 10.61A a single loop and in Figure 10.61D a single tap connects the cavity across or shunts the transmission line. At resonance, the open ¼-λ cavity returns energy to the transmission line through the loop, canceling energy in the line.

BP-BR CAVITIES

A third class of cavity filter has both BP and BR characteristics. In Figures 10.61B and C the cavity is connected across a coil or capacitor. All BR cavities using a shunt inductor or capacitor also have a band-pass response "bump" as shown in Figure 10.63B. The combination of the cavity's reactance (shown in Figure 10.62) in parallel with the reactance inserted into the transmission line by the inductor or capacitor create a high-Q resonance that blocks the energy in the line at the notch frequency which is offset from the cavity's natural resonance.

If a shunt inductor is used, as in Figure 10.61B, the notch will be above the cavity's natural resonance where its capacitive reactance resonates with the shunt inductive reactance. If a shunt capacitor is used, as in Figure 10.61C, the notch will be on the low frequency side. The amount of inductance or capacitance determines the separation of the notch and bump in the cavity's response. (If the band-pass bump is intentionally minimized, the cavity is called a *pure notch cavity*.)

To be used in a repeater duplexer, it is possible to configure the bump frequency to be above or below that of the notch. As in Figure 10.63B, if the transmitter frequency (the single to be "notched out") is higher than the receiver frequency, such as for a 2 meter repeater with a transmit frequency of 146.97 MHz and a receive frequency of 146.37 MHz, the bump (the band-pass range) should be below the notch and vice versa. The bump is typically adjusted to be centered on the receive frequency. (The use of cavities for repeater duplexers is discussed in more detail in the **Repeater Systems** chapter.)

For more detailed discussions of cavity operation and photos of a disassembled commercial cavity showing several of the techniques discussed here, see **www.repeater-builder.com/antenna/pdf/ve2azx-duplexer-info.pdf** or the online article by ZL1BTB at **www.amalgamate2000.com/radio-hobbies/radio/coaxial_resonators.htm**. The Repeater-Builders web page **www.repeater-builder.com/antenna/ant-sys-index.html#duplexers** offers a number of excellent articles, as well.

Building high-quality microwave cavities is illustrated in the *QEX* articles "Easy Microwave Filters Using Waveguides and Cavities" and "Very High Q Microwave Cavities and Filters" by Paolo Antoniazzi, IW2ACD, which are included with the online content.

10.9.4 Helical Resonators

Ever-increasing occupancy of the radio spectrum brings with it a parade of receiver overload and spurious responses. Overload problems can be minimized by using high-dynamic-range receiving techniques, but

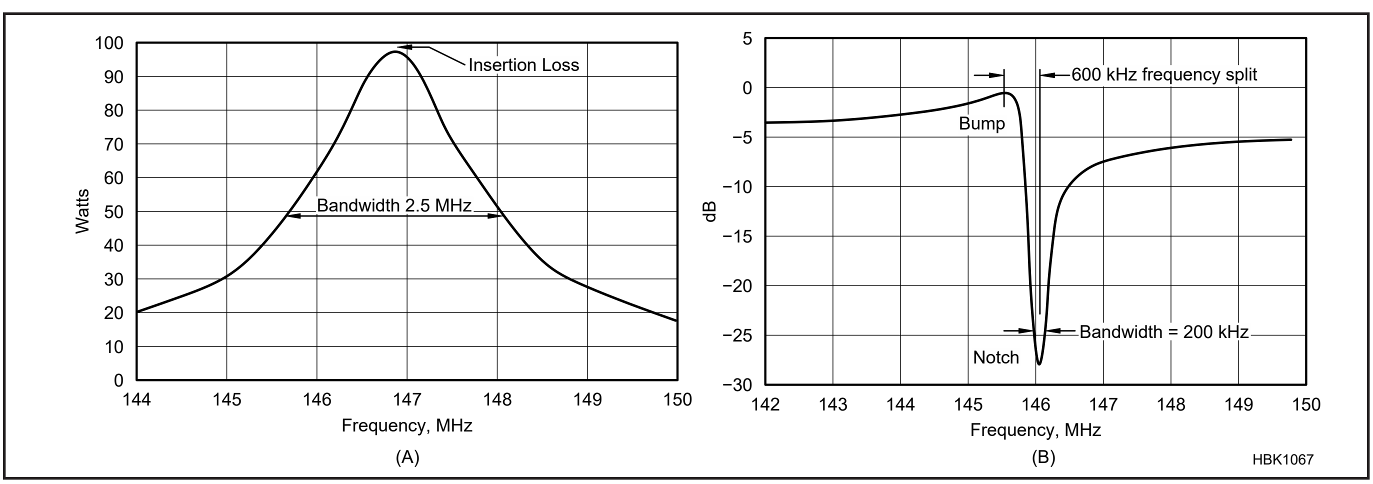

Figure 10.63 — Typical band-pass (BP) and notch (BR) response for 2 meter cavity filters.

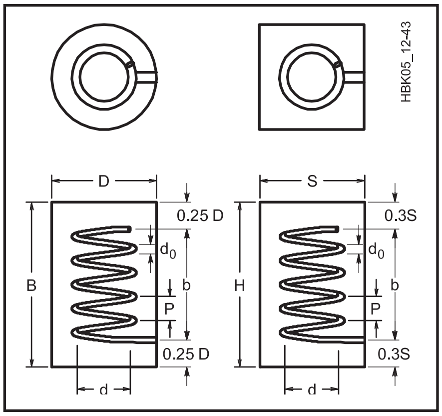

Figure 10.64 — Dimensions of round and square helical resonators. The diameter, D (or side, S) is determined by the desired unloaded Q. Other dimensions are expressed in terms of D or S (see text).

spurious responses (such as the image frequency) must be filtered out before mixing occurs. Conventional tuned circuits cannot provide the selectivity necessary to eliminate the plethora of signals found in most urban and many suburban environments. Other filtering techniques must be used.

Helical resonators are usually a better choice than ¼-λ cavities on 50, 144 and 222 MHz to eliminate these unwanted inputs because they are smaller and easier to build than coaxial cavity resonators, although their Q is not as high as that of cavities. In the frequency range from 30 to 100 MHz it is difficult to build high-Q inductors, and coaxial cavities are very large. In this frequency range the helical resonator is an excellent choice. At 50 MHz for example, a capacitively tuned, ¼-λ coaxial cavity with an unloaded Q of 3000 would be about 4 inches in diameter and nearly 5 ft long. On the other hand, a helical resonator with the same unloaded Q is about 8.5 inches in diameter and 11.3 inches long. Even at 432 MHz, where coaxial cavities are common, the use of helical resonators results in substantial size reductions.

The helical resonator was described by the late Jim Fisk, W1HR, in a June 1976 *QST* article. (see the **Reference** section) The resonator is described as a coil surrounded by a shield, but it is actually a shielded, resonant section of helically wound transmission line with relatively high characteristic impedance and low propagation velocity along the axis of the helix. The electrical length is about 94% of an axial ¼ -λ or 84.6°. One lead of the helical winding is connected directly to the shield and the other end is open circuited as shown in **Figure 10.64**. Although the shield may be any shape, only round and square shields will be considered here.

10.9.5 Helical Resonator Design

The unloaded Q of a helical resonator is determined primarily by the size of the shield. For a round resonator with a copper coil on a low-loss form, mounted in a copper shield, the unloaded Q is given by

$$Q_U = 50D\sqrt{f_0}$$

where

D = inside diameter of the shield, in inches, and

f_0 = frequency, in MHz.

D is assumed to be 1.2 times the width of one side for square shield cans. This formula includes the effects of losses and imperfections in practical materials. It yields values of unloaded Q that are easily attained in practice. Silver plating the shield and coil increases the unloaded Q by about 3% over that predicted

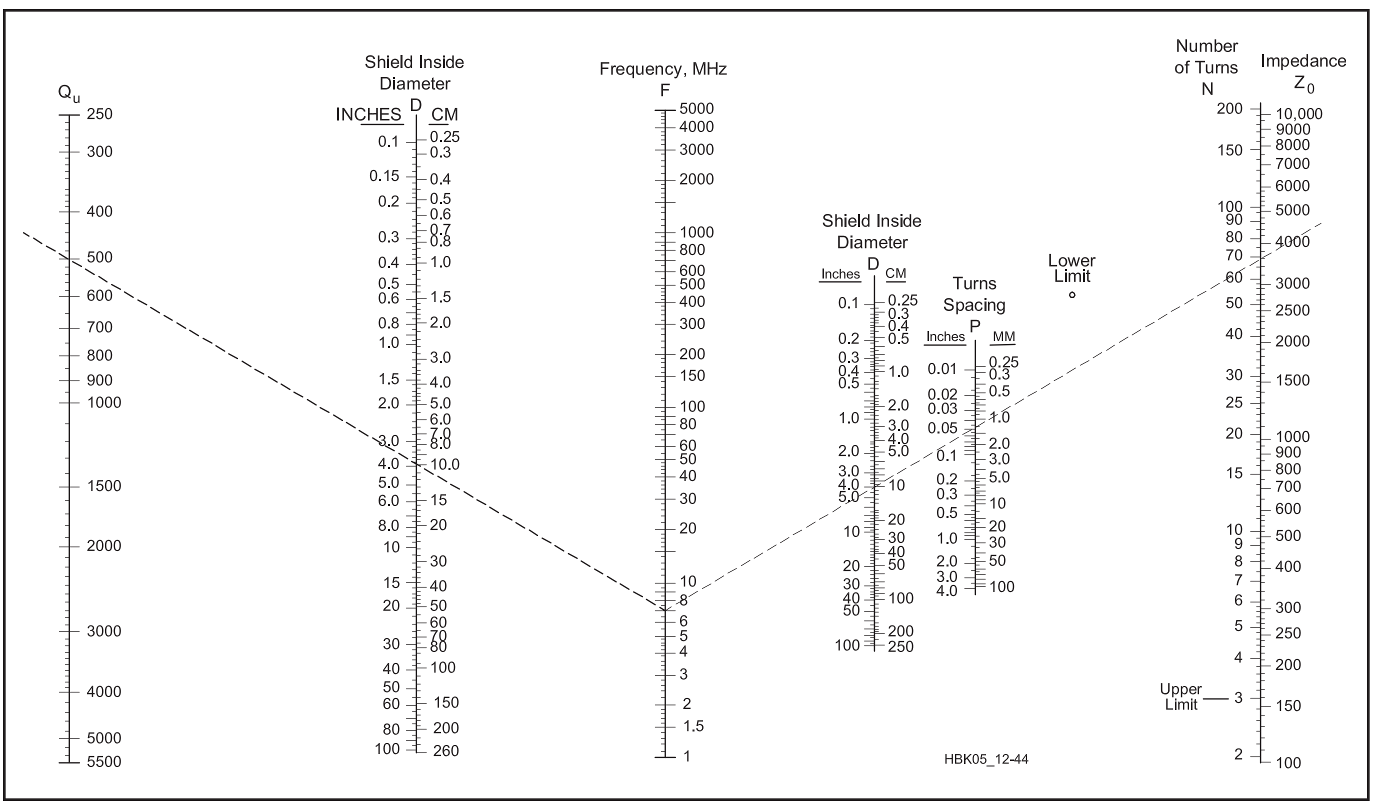

Figure 10.65 — The design nomograph for round helical resonators starts by selecting Q_U and the required shield diameter. A line is drawn connecting these two values and extended to the frequency scale (example here is for a shield of about 3.8 inches and Q_U of 500 at 7 MHz). Finally the number of turns, N, winding pitch, P, and characteristic impedance, Z_0, are determined by drawing a line from the frequency scale through selected shield diameter (but this time to the scale on the right-hand side. For the example shown, the dashed line shows P ≈ 0.047 inch, N = 70 turns, and Z_n = 3600 Ω).

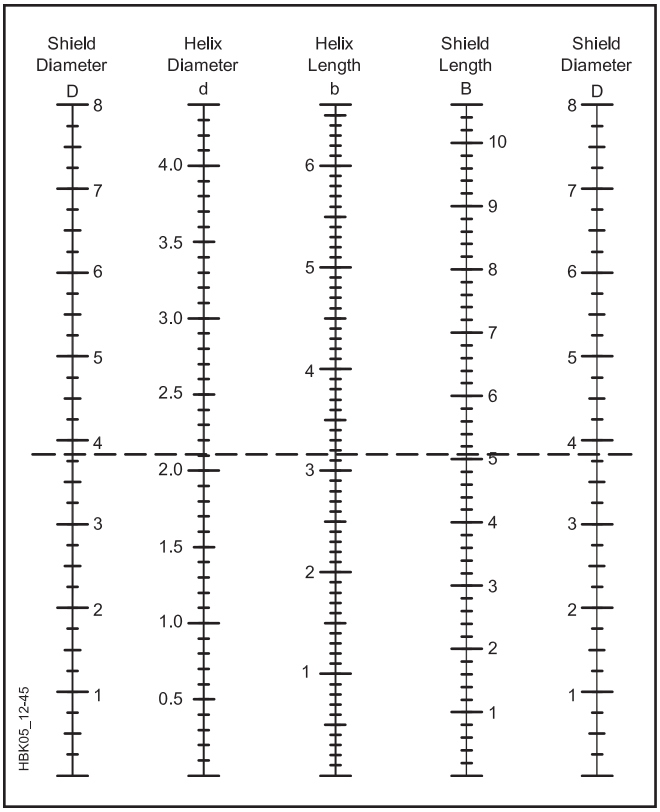

Figure 10.66 — The helical resonator is scaled from this design nomograph. Starting with the shield diameter, the helix diameter, d, helix length, b, and shield length, B, can be determined with this graph. The example shown has a shield diameter of 3.8 inches. This requires a helix mean diameter of 2.1 inches, helix length of 3.1 inches, and shield length of 5 inches.

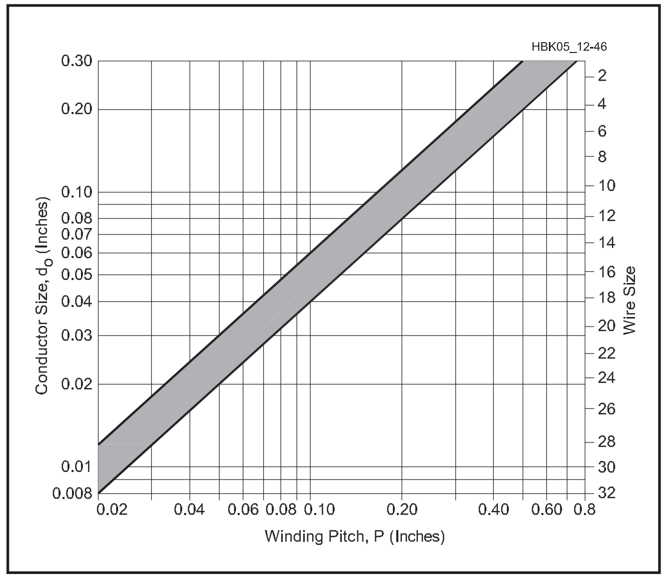

Figure 10.67 — This chart provides the design information of helix conductor size vs winding pitch, P. For example, a winding pitch of 0.047 inch results in a conductor diameter between 0.019 and 0.028 inch (#22 or #24 AWG).

by the equation. At VHF and UHF, however, it is more practical to increase the shield size slightly (that is, increase the selected QU by about 3% before making the calculation). The fringing capacitance at the open-circuit end of the helix is about 0.15D pF (that is, approximately 0.3 pF for a shield 2 inches in diameter). Once the required shield size has been determined, the total number of turns, N, winding pitch, P and characteristic impedance, Z_0, for round and square helical resonators with air dielectric between the helix and shield, are given by:

$$N = \frac{1908}{f_0 D}$$

$$P = \frac{f_0 D^2}{2312}$$

$$Z_0 = \frac{99,000}{f_0 D}$$

$$N = \frac{1590}{f_0 S}$$

$$P = \frac{f_0 S^2}{1606}$$

$$Z_0 = \frac{82,500}{f_0 S}$$

In these equations, dimensions D and S are in inches and f_0 is in megahertz. The design nomograph for round helical resonators in **Figure 10.65** is based on these formulas.

Although there are many variables to consider when designing helical resonators, certain ratios of shield size to length and coil diameter to length, provide optimum results. For helix diameter, d = 0.55 D or d = 0.66 S. For helix length, b = 0.825D or b = 0.99S. For shield length, B = 1.325 D and H = 1.60 S.

Design of filter dimensions can be done using the nomographs in this section or with computer software. The program *Helical* for designing and analyzing helical filters is available with the online content. Use of the nomographs is described in the following paragraphs.

Figure 10.66 simplifies calculation of these dimensions. Note that these ratios result in a helix with a length 1.5 times its diameter, the condition for maximum Q. The shield is about 60% longer than the helix — although it can be made longer — to completely contain the electric field at the top of the helix and the magnetic field at the bottom.

The winding pitch, P, is used primarily to determine the required conductor size. Adjust the length of the coil to that given by the equations during construction. Conductor size ranges from 0.4 P to 0.6 P for both round and square resonators and are plotted graphically in **Figure 10.67**.

Obviously, an area exists (in terms of frequency and unloaded Q) where the designer must make a choice between a conventional cavity (or lumped LC circuit) and a helical resonator. The choice is affected by physical shape at higher frequencies. Cavities are long and relatively small in diameter, while the length of a helical resonator is not much greater than its diameter. A second consideration is that point where the winding pitch, P, is less than the radius of the helix (otherwise the structure tends to be non-helical). This condition occurs when the helix has fewer than three turns (the "upper limit" on the design nomograph of Figure 10.65).

10.9.6 Helical Filter Construction

The shield should not have any seams parallel to the helix axis to obtain as high an unloaded Q as possible. This is usually not a problem with round resonators because large-diameter copper tubing is used for the shield, but square resonators require at least one seam and usually more. The effect on unloaded Q is minimized if the seam is silver soldered carefully from one end to the other.

Results are best when little or no dielectric is used inside the shield. This is usually no problem at VHF and UHF because the con-

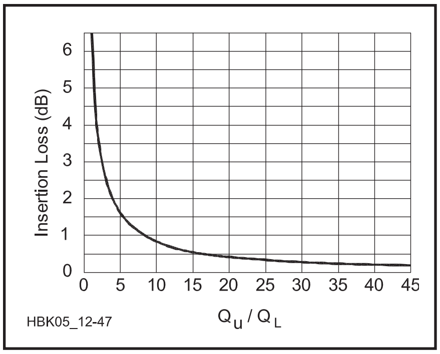

Figure 10.68 — The ratio of loaded (Q_L) to unloaded (Q_U) Q determines the insertion loss of a tuned resonant circuit.

ductors are large enough that a supporting coil form is not required. The lower end of the helix should be soldered to the nearest point on the inside of the shield.

Although the external field is minimized by the use of top and bottom shield covers, the top and bottom of the shield may be left open with negligible effect on frequency or unloaded Q. Covers, if provided, should make electrical contact with the shield. In those resonators where the helix is connected to the bottom cover, that cover must be soldered solidly to the shield to minimize losses.

10.9.7 Helical Resonator Tuning

A carefully built helical resonator designed from the nomograph of Figure 10.65 will resonate very close to the design frequency. Slightly compress or expand the helix to adjust resonance over a small range. If the helix is made slightly longer than that called for in Figure 10.66, the resonator can be tuned by pruning the open end of the coil. However, neither of these methods is recommended for wide frequency excursions because any major deviation in helix length will degrade the unloaded Q of the resonator.

Most helical resonators are tuned by means of a brass tuning screw or high-quality air-variable capacitor across the open end of the helix. Piston capacitors also work well, but the Q of the tuning capacitor should ideally be several times the unloaded Q of the resonator. Varactor diodes have sometimes been used

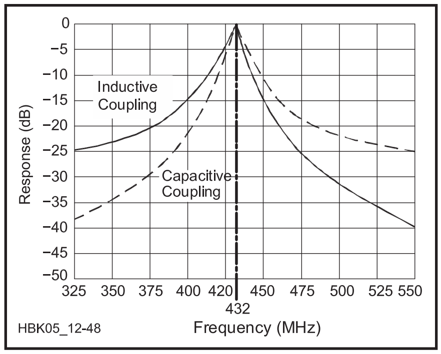

Figure 10.69 — This response curve for a single-resonator 432 MHz filter shows the effects of capacitive and inductive input/output coupling. The response curve can be made symmetrical on each side of resonance by combining the two methods (inductive input and capacitive output, or vice versa).

where remote tuning is required, but varactors can generate unwanted harmonics and other spurious signals if they are excited by strong, nearby signals.

When a helical resonator is to be tuned by a variable capacitor, the shield size is based on the chosen unloaded Q at the operating frequency. Then the number of turns, N and the winding pitch, P, are based on resonance at 1.5 f_0. Tune the resonator to the desired operating frequency, f_0.

10.9.8 Helical Resonator Insertion Loss

The insertion loss (dissipation loss), I_L, in decibels, of all single-resonator circuits is given by

$$I_L = 20 \log_{10} \left(\frac{1}{1 - \frac{Q_L}{Q_U}} \right)$$

where

Q_L = loaded Q, and

Q_U = unloaded Q.

This is plotted in **Figure 10.68**. For the most practical cases ($Q_L > 5$), this can be closely approximated by $I_L \approx 9.0 \, (Q_L/Q_U)$ dB. The selection of Q_L for a tuned circuit is dictated primarily by the required selectivity of the circuit. However, to keep dissipation loss to 0.5 dB or less (as is the case for low-noise VHF receivers), the unloaded Q must be at least 18 times the Q_L.

10.9.9 Coupling Helical Resonators

Signals are coupled into and out of helical resonators with inductive loops at the bottom of the helix, direct taps on the coil or a combination of both. Although the correct tap point can be calculated easily, coupling by loops and probes must be determined experimentally.

The input and output coupling is often provided by probes when only one resonator is used. The probes are positioned on opposite sides of the resonator for maximum isolation. When coupling loops are used, the plane of the loop should be perpendicular to the axis of the helix and separated a small distance from the bottom of the coil. For resonators with only a few turns, the plane of the loop can be tilted slightly so it is parallel with the slope of the adjacent conductor.

Helical resonators with inductive coupling (loops) exhibit more attenuation to signals above the resonant frequency (as compared to attenuation below resonance), whereas resonators with capacitive coupling (probes) exhibit more attenuation below the pass-band, as shown for a typical 432 MHz resonator in **Figure 10.69**. Consider this characteristic when choosing a coupling method. The pass-band can be made more symmetrical by using a combination of coupling methods (inductive input and capacitive output, for example).

If more than one helical resonator is required to obtain a desired band-pass characteristic, adjacent resonators may be coupled through apertures in the shield wall between the two resonators. Unfortunately, the size and location of the aperture must be found empirically, so this method of coupling is not very practical unless you're building a large number of identical units.

Since the loaded Q of a resonator is determined by the external loading, this must be considered when selecting a tap (or position of a loop or probe). The ratio of this external loading, R_b, to the characteristic impedance, Z_0, for a ¼-λ resonator is calculated from:

$$K = \frac{R_b}{Z_0} = 0.785 \left(\frac{1}{Q_L} - \frac{1}{Q_U} \right)$$

10.10 HF Transmitting Filters

Multi-transmitter stations operate transmitters and receivers in close proximity, either for contesting, special events, DXpeditions and Field Day operation, or emergency and disaster response communications. In such stations, band-pass filters (BPFs) are commonly used to reduce interstation interference from harmonics and broadband noise. Band-pass filters also make receivers less susceptible to overload from nearby transmitters. Remember that when trying to eliminate harmonics or noise, the filter must be installed at the *source*. Once transmitted and radiated, these spurious signals act as in-band signals at the receiver and cannot be filtered out.

This section is based on material from articles by Ethan Miller, K8GU; Edward Wetherhold, W3NQN; and Jeffrey Crawford, KØZR, referenced at the end of this section. The most popular HF BPF designs in use today are the 200 W designs by W3NQN. They can be built from individual components, as kits, and as complete units from several manufacturers. Complete design and building information are available in the two *QST* articles by W3NQN that are referenced at the end of this section. Another set of designs by James Tonne, W4ENE, for lower-power applications to meet the FCC emissions limits in §97.307 is provided in the Projects section of this chapter.

In addition, two publications are important complements to this section. The first is *Managing Interstation Interference* by George Cutsogeorge, W2VJN, available at no cost from Inrad at **www.inrad.net/home.php**. The other is *Low-Band DXing* by John Devoldere, ON4UN. ON4UN's chapter on contest station engineering contains good information on filtering techniques.

10.10.1 Transmitting Filter Basics

There are two places in the station to install transmitting filters: Between the transceiver and amplifier and after the amplifier at the input to the antenna system. They operate in both the receive path and transmit paths (unless separate receive antennas are used). These filters reduce the level of transmitted harmonics and broadband noise and help to protect the receiver from nearby transmitters.

Transmitting filters can be constructed from transmission line stubs or discrete components. At high power levels stubs tend to be less expensive than discrete component band-pass filters. They also have the advantage that they can be designed to work well at harmonically related frequencies.

The two primary performance requirements for transmitting filters are for minimum insertion loss and maximum attenuation at the second harmonic, which is usually the strongest and most troublesome harmonic in a multi-station environment along with wideband noise from the transmitter's oscillators and synthesizers.

In practice, filters built from discrete components are usually capable of greater attenuation than stubs although the power-handling capability of a stub-based filter may make it a more appropriate choice. For exciter power levels (i.e., <200 W), discrete filters tend to be quite economical. However, if the transceiver or antenna is not well matched to the filter, the effectiveness of the system is reduced. This situation is very much installation specific and difficult to evaluate in advance.

Insertion loss for transmitting filters should be significantly less than 1.0 dB which is equivalent to approximately 20 percent of the transmitted power. It is recommended to use a vector network analyzer (VNA — see the **Test Instruments and Measurements** chapter) when tuning or adjusting the filters. It is important to measure SWR into the filter, loss through the filter, and attenuation on adjacent bands. Measurements must be made with the filter cover in place to include the effects of component-to-enclosure proximity.

It is important that transmitting filters be operated with a low SWR at the output either looking into an amplifier input or into the antenna system. The filter capacitors are usually the most susceptible to failure by overheating when transmitting into a high SWR load. This can occur when transmitting into the "wrong" antenna or transmitting into the filter on the "wrong" band. An automatic antenna tuner can also cause filter failure by allowing full power to be input to the filter, even if on the wrong band or antenna. If possible, use the transceiver (and amplifier) band output data to automatically switch equipment and antennas to the "right" setting.

BUILDING TRANSMITTING FILTERS

Transmitting filters require sturdy components that can handle high current and high voltage, which increases their size and cost. At and above 100 W, toroidal inductors are generally used because of their self-shielding properties that allows the inductors to be physically close with little interaction, making for a more compact filter. As an example, **Figure 10.70** shows a W3NQN filter constructed by K8GU.

For power levels below 5 W, toroid core sizes of 0.5 inches outer diameter or less are commonly used. At power levels of 150 to 200 W, however, the core size must be much larger to dissipate the heat resulting from core and winding losses without excessive temperature rise. It was once assumed that core saturation was the primary limiting factor in high-power RF applications. However, excessive temperature increases resulting from losses in the winding and core material is the limiting factor. A temperature increase of less than 40 °C is preferred so that in an ambient of 90 °F (typical temperature for a hot day), the core temperature will be not more than 32 + 40 °C = 72 °C, or well below 90 °C.

To minimize time spent adjusting inductor values in the assembled filter, use an LC meter (see the **Test Equipment and Measurements** chapter) to adjust the inductor to the exact value before soldering it into the final assembly. Capacitors should be measured, as well. Capacitors can be combined in parallel to reach a final, exact value. This has the additional benefit of reducing current and heating in the capacitors.

Most transmitting filters at the 100 W level use high-voltage disc ceramic capacitors. To minimize the chance of failure caused by greater-than-anticipated voltages or currents, and to derate the standard dc voltage rating by 50% for RF applications, capacitors with dc ratings of 2, 3, and 4 kV are used. To minimize

Figure 10.70 — K8GU band-pass Filter.

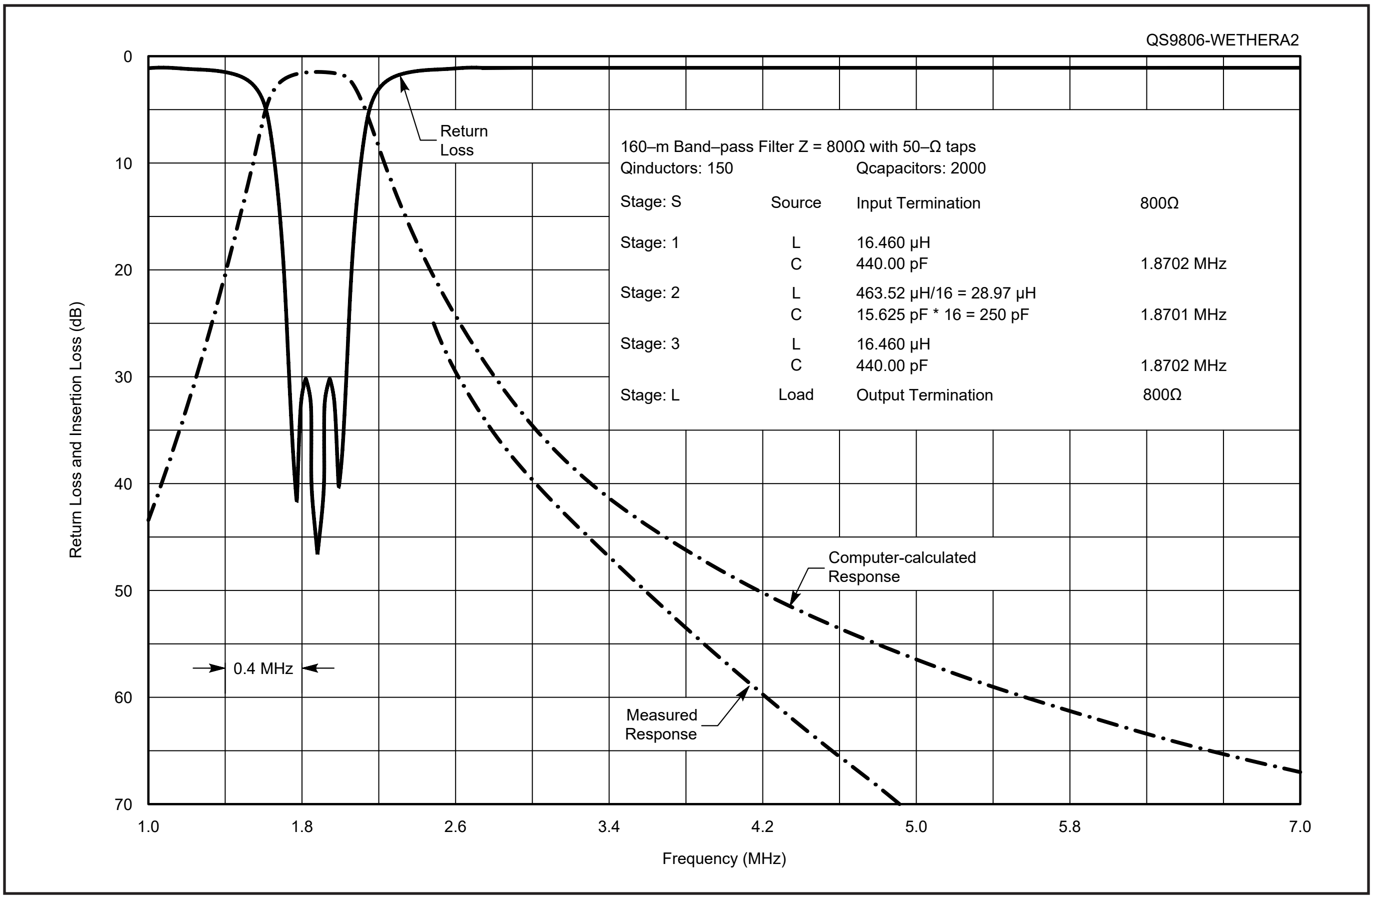

Figure 10.71 — Predicted and as-built performance of a W3NQN 160-meter BPF.

the temperature rise caused by several amps of current, two or more capacitors in parallel can be used. For high currents, CDV-16 capacitors from CDE are rated for 5 A of continuous current at HF.

Adjustable mica compression trimmers (usually made by Arco) were popular components for tunable filters but are no longer available except as surplus and new old stock (NOS). Trimmer capacitors available now are generally low-value (<50 pF, max.) and not rated above 200 V. High-quality Teflon or semi-rigid coax can also be substituted for low-value capacitors and trimmed to the proper value as described in the article "A Coaxial Alternative to low value SBH capacitors" by Bob Henderson, 5B4AGN, at **groups.io/g/TXBPF/files/A%20Coaxial%20Alternative%20to%20low%20value%20SBH%20capacitors.pdf**.

Once the filter has been constructed, the easiest way to complete tuning is to use either a spectrum analyzer with tracking generator or a VNA. Compare the actual behavior to the design's predicted value as in **Figure 10.71** which shows a W3NQN 160 meter BPF. The filter designer may also provide tuning instructions.

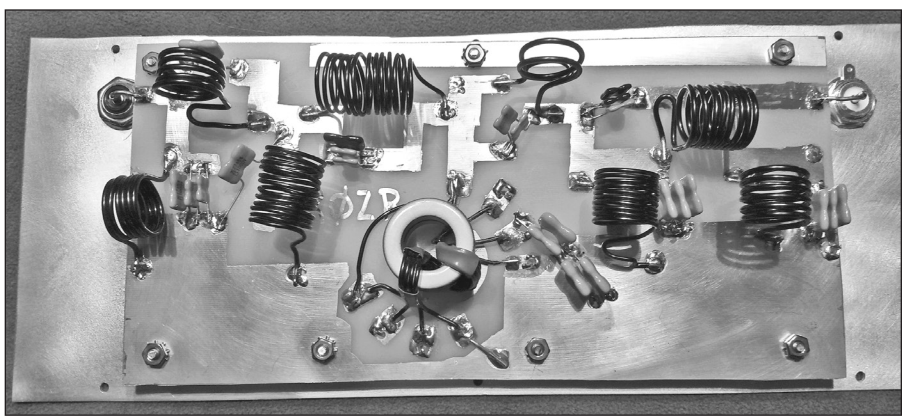

Figure 10.72 — A high-power 20 meter BPF assembly by KØZR.

10.10.2 High-Power Transmitting Filters

It is important in a high-power filter design to design carefully and leave little room for error. For example, consider a filter with a dissipative insertion loss of 0.3 dB. With 1500 W applied, 100 W will be dissipated in the filter. This heat will age the components more rapidly, may shift the filter's return loss as a function of duty cycle, and may unnecessarily cause thermal stresses to the filter, leading to possible premature failure. Heat is one of the enemies in high-power filter design.

The voltages and currents encountered in a high-power filter can be very high, especially when the filter is operated with loads far from 50 Ω. Circuit-analysis tools that show conditions for individual components should be used to perform "what if" calculations such as, "what happens to RF currents if the SWR were to be 2:1 instead of 1:1?"

Current-carrying inductors should be wound with #12 or #14 AWG Thermaleze-coated wire. Capacitors can be placed in series when higher breakdown voltages are required, and similarly, placed in parallel to increase the net current capacity of a given capacitor. The article on high-power filters by W2VJN presents a technique of using shorted coaxial stubs. **Figure 10.72** shows a high-power 20 meter BPF constructed by KØZR. Note that is mounted on a metal plate with the input and output SO-239 connectors at the left and right.

Because of the extra heat dissipation in filters used at high power, it is also important to consider ventilation and cooling. Enclosures should have holes or slots that are big enough to allow for air circulation. After assembling a filter for the first time, it is recommended to do a heat-soak test at full power and measure the temperature of all current-carrying components.

10.10.3 Transmitting Filter References

Baker, A., KG4JJH, "Contest Band-pass Filters," **www.kg4jjh.com/pdf/Contest%20Band-pass%20Filters.pdf**.

Brown, J., K9YC, "Band-pass Filter Survey," Jim Brown, **audiosystemsgroup.com/BandpassFilterSurvey.pdf**.

Crawford, J., KØZR, "High-Power HF Band-Pass Filter Design," *QEX*, Mar. 2018, pp. 3 – 10.

Cutsogeorge, G., W2VJN, "High-Power Harmonic Filters," *NCJ*, Nov. 1998, pp.14 – 15.

Miller, E., K8GU, "Selecting Band-Pass Filters and Switching Hardware," *NCJ*, Mar. 2008, pp. 29 – 31. Additional information at **www.k8gu.com/engineering/w3nqn-tx-filters**.

Wetherhold, E., W3NQN, "Clean Up Your Signals with Band-Pass Filters," *QST*, Part 1 May 1998; Part 2 Jun. 1998.

10.10.4 Diplexer Filter

This section, covering diplexer filters, was written by William E. Sabin, WØIYH. The diplexer is helpful in certain applications, such as frequency mixer terminations.

The terms "diplexer" and "duplexer" are often confused. A duplexer is a device that allows transmitters and receivers to share a common antenna on a single band. A duplexer can be a device such as a circulator or isolator (types of waveguide components that allow power to flow in only one direction) or filters (cavity or lumped-circuit). Diplexers are filters that enable antennas on different bands to be used by a transceiver through a common feed line. Diplexers can also be used to allow transceivers on different bands to share a multiband antenna. A triplexer extends the diplexer to three bands.

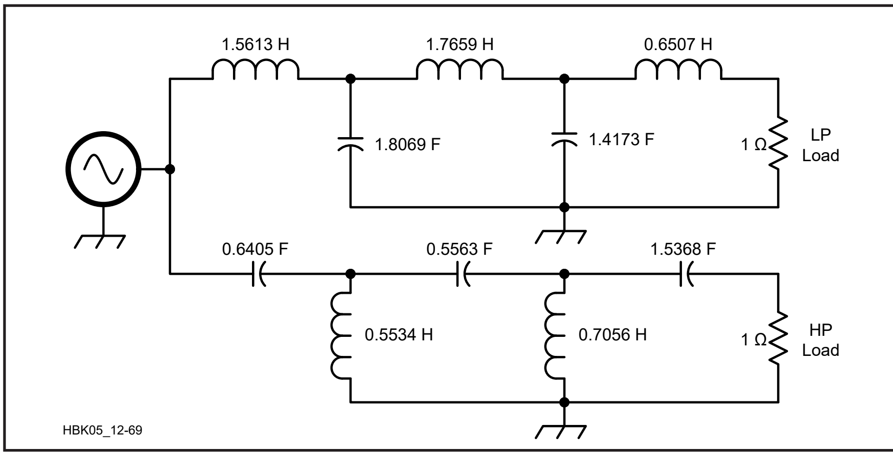

Figure 10.73 — Low-pass and high-pass prototype diplexer filter design. The low-pass portion is at the top, and the high-pass at the bottom of the drawing. See text.

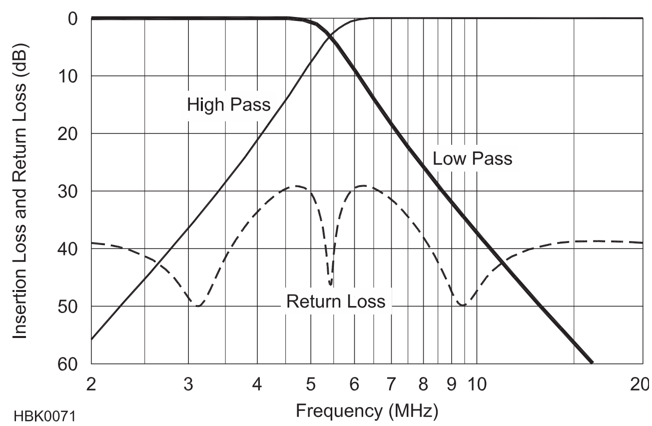

Figure 10.74 — Response for the low-pass and high-pass portions of the 80 meter diplexer filter. Also shown is the return loss of the filter.

It is important that a diplexer or triplexer have a constant filter-input resistance that extends to the stop band as well as the passband. Ordinary filters that become highly reactive or have an open or short-circuit input impedance outside the passband may degrade performance of the devices to which they are attached. For example, impedances far from 50 Ω outside the operating frequency range may cause an amplifier to become unstable and generate spurious emissions or oscillate.

Figure 10.73 shows a *normalized* prototype 5-element, 0.1-dB Chebyshev low-pass/high-pass (LP/HP) filter. This idealized filter is driven by a voltage generator with zero internal resistance, has load resistors of 1.0 Ω and a cutoff frequency of 1.0 radian per second (0.1592 Hz). The LP prototype values are taken from standard filter tables.[1] The first element is a series inductor. The HP prototype is found by:

a) replacing the series L (LP) with a series C (HP) whose value is 1/L, and

b) replacing the shunt C (LP) with a shunt L (HP) whose value is 1/C.

For the Chebyshev filter, the return loss is improved several dB by multiplying the prototype LP values by an experimentally derived number, K, and dividing the HP values by the same K. You can calculate the LP values in henrys and farads for a 50 Ω RF application with the following formulas:

$$L_{LP} = \frac{KL_{P(LP)}R}{2\pi f_{CO}}; C_{LP} = \frac{KC_{P(LP)}}{2\pi f_{CO}R}$$

where

$L_{P(LP)}$ and $C_{P(LP)}$ are LP prototype values,
K = 1.005 (in this specific example),
R = 50 W, and
f_{CO} = the cutoff (–3 dB response) frequency in Hz.

For the HP segment:

$$L_{HP} = \frac{L_{P(HP)}R}{2\pi f_{CO}K}; C_{HP} = \frac{C_{P(HP)}}{2\pi f_{CO}KR}$$

where $L_{P(HP)}$ and $C_{P(HP)}$ are HP prototype values.

Figure 10.74 shows the LP and HP responses of a diplexer filter for the 80 meter band. The following items are to be noted:

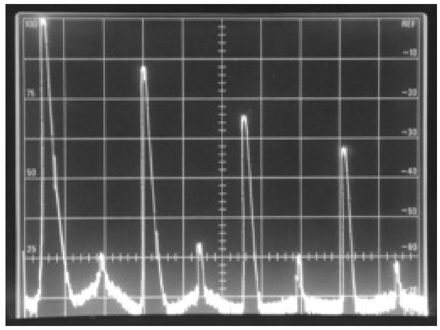

(A)

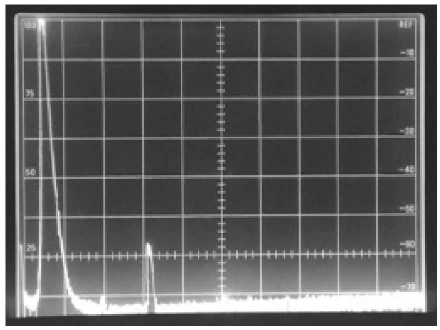

(B)

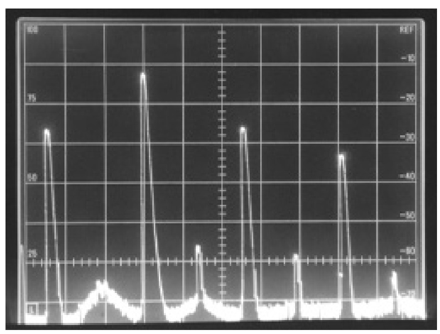

(C)

Figure 10.75 — At A, the output spectrum of a push-pull 80 meter amplifier. At B, the spectrum after passing through the low-pass filter. At C, the spectrum after passing through the high-pass filter.

• The 3 dB responses of the LP and HP meet at 5.45 MHz.

• The input impedance is close to 50 Ω at all frequencies, as indicated by the high value of return loss (SWR <1.07:1).

• At and near 5.45 MHz, the LP input reactance and the HP input reactance are conjugates; therefore, they cancel and produce an almost perfect 50 Ω input resistance in that region.

• Because of the way the diplexer filter is derived from synthesis procedures, the transfer characteristic of the filter is mostly independent of the actual value of the amplifier dynamic output impedance.2 This is a useful feature, since the RF power amplifier output impedance is usually not known or specified.

• The 80 meter band is well within the LP response.

• The HP response is down more than 20 dB at 4 MHz.

• The second harmonic of 3.5 MHz is down only 18 dB at 7.0 MHz. Because the second harmonic attenuation of the LP is not great, it is necessary that the amplifier itself be a well-balanced push-pull design that greatly rejects the second harmonic. In practice this is not a difficult task.

• The third harmonic of 3.5 MHz is down almost 40 dB at 10.5 MHz.

Figure 10.75A shows the unfiltered output of a solid-state push-pull power amplifier for the 80 meter band. In the figure you can see that:

• The second harmonic has been suppressed by a proper push-pull design.

• The third harmonic is typically only 15 dB or less below the fundamental.

The amplifier output goes through our diplexer filter. The desired output comes from the LP side, and is shown in Figure 10.75B. In it we see that:

• The fundamental is attenuated only about 0.2 dB.

• The LP has some harmonic content; however, the attenuation exceeds FCC requirements for a 100 W amplifier.

Figure 10.75C shows the HP output of the diplexer that terminates in the HP load or dump resistor. A small amount of the fundamental frequency (about 1%) is also lost in this resistor. Within the 3.5 to 4.0 MHz band, the filter input resistance is almost exactly the correct 50 W load resistance value. This is because power that would otherwise be reflected back to the amplifier is absorbed in the dump resistor.

Solid state power amplifiers tend to have stability problems that can be difficult to debug.3 These problems may be evidenced by level changes in: load impedance, drive, gate or base bias, supply voltage, etc. Problems may arise from:

• The reactance of the low-pass filter outside the desired passband. This is especially true for transistors that are designed for high-frequency operation.

• Self-resonance of a series inductor at some high frequency.

• A stop band impedance that causes voltage, current and impedance reflections back to the amplifier, creating instabilities within the output transistors.

Intermodulation performance can also be degraded by these reflections. The strong third harmonic is especially bothersome for these problems.

The diplexer filter is an approach that can greatly simplify the design process, especially for the amateur with limited PA-design experience and with limited home-lab facilities. For these reasons, the amateur homebrew enthusiast may want to consider this solution, despite its slightly greater parts count and expense.

The diplexer is a good technique for narrowband applications such as the HF amateur bands.[4] From Figure 10.74, we see that if the signal frequency is moved beyond 4.0 MHz the amount of desired signal lost in the dump resistor becomes large. For signal frequencies below 3.5 MHz the harmonic reduction may be inadequate. A single filter will not suffice for all the HF amateur bands.

This treatment provides you with the information to calculate your own filters. A *QEX* article has detailed instructions for building and testing a set of six filters for a 120 W amplifier. These filters cover all nine of the MF/HF amateur bands.[5]

You can use this technique for other filters such as Bessel, Butterworth, linear phase, Chebyshev 0.5, 1.0, etc.[6] However, the diplexer idea does *not* apply to the elliptic function types.

The diplexer approach is a resource that can be used in any application where a constant value of filter input resistance over a wide range of passband and stop band frequencies is desirable for some reason. Computer modeling is an ideal way to finalize the design before the actual construction. The coil dimensions and the dump resistor wattage need to be determined from a consideration of the power levels involved.

Another significant application of the diplexer is for elimination of EMI, RFI and TVI energy. Instead of being reflected and very possibly escaping by some other route, the unwanted energy is dissipated in the dump resistor.[7]

See the discussion "Design Software for LC Filters" at the end of the Passive LC Filters section of this chapter. The software package provided by Jim Tonne, W4ENE, of Tonne Software (**www.tonnesoftware.com**) includes *Diplexer* which greatly simplifies the process of designing diplexer filters. The software is part of the online content for this book.

Notes

[1]Williams, A. and Taylor, F., *Electronic Filter Design Handbook*, any edition, McGraw-Hill.

[2]Storer, J., *Passive Network Synthesis* (McGraw-Hill, 1957), pp. 168 – 170. This book shows that the input resistance is ideally constant in the passband and the stop band and that the filter transfer characteristic is ideally independent of the generator impedance.

[3]Sabin, W. and Schoenike, E., *HF Radio Systems and Circuits* (Noble Publishing, 1998).

[4]Dye, N. and Granberg, H., *Radio Frequency Transistors, Principles and Applications* (Butterworth-Heinemann, 1993) p. 151.

[5]Sabin, W., WØIYH, "Diplexer Filters for the HF MOSFET Power Amplifier," *QEX*, July/Aug. 1999. Also check the ARRL website at **www.arrl.org/qex**.

[6]See note 1. *Electronic Filter Design Handbook* has LP prototype values for various filter types, and for complexities from 2 to 10 components.

[7]Weinrich, R. and Carroll, R., "Absorptive Filters for TV Harmonics," *QST*, Nov. 1968, pp. 10 – 25.

10.11 Filter Projects

The filter projects to follow are by no means the only filter projects in this book. Filters for specific applications may be found in other chapters of this *Handbook* and in its online content. Receiver input filters, transmitter filters, inter-stage filters and others can be extracted from the various projects and built for other applications. Since filters are a first line of defense against *electromagnetic interference* (EMI) problems, additional filter methods appear in the **RF Interference** chapter.

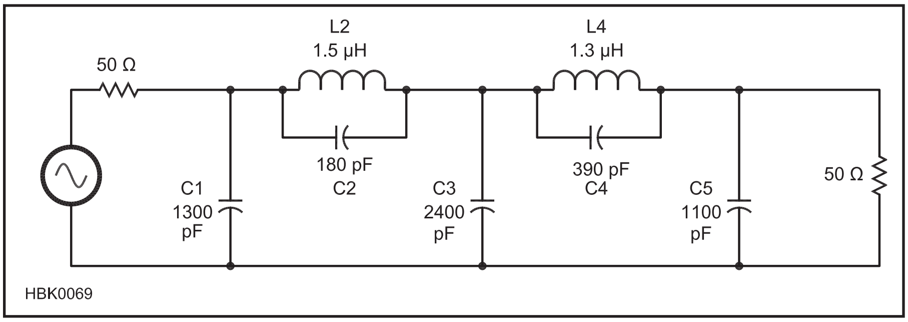

Figure 10.76 — Optimized low-pass filter. This design is for the 80 meter amateur band. It is similar to a Cauer design but the parts values have been optimized as described in the text and in the Sep/Oct 1998 issue of *QEX*.

10.11.1 Optimized Harmonic Transmitting Filters

Low-pass filters should be placed at the output of transmitters to ensure that they meet the various regulatory agency requirements for harmonic suppression. These are commonly designed to pass a single amateur band and provide attenuation at harmonics of that band sufficient to meet the requirements. The material presented here by Jim Tonne, W4ENE, is based on material originally published in the September/October 1998 issue of *QEX*. Some 100-W and higher power band-pass filters are discussed in articles referenced by this chapter's previous section HF Transmitting Filters which are included in the online content.

The basic approach is to use a computer to optimize the performance in the passband (a single amateur band) while simultaneously maximizing the attenuation at the second and third harmonic of that same band. When this is done, the higher harmonics will also be well within spec.

The schematic of this filter along with parts values for the 3.5 to 4.0 MHz amateur band is shown in **Figure 10.76**. The responses of that filter are shown in **Figure 10.77**.

Component values for the 160 meter through the 6 meter amateur bands are shown in **Table 10.5**. The capacitors are shown in pF and the inductors in µH. The capacitors are the nearest 5% values; both the nearest 5% and the exact inductor values are shown.

Using the nearest-5% inductor values will result in satisfactory operation. If the construction method is such that exact-value (adjustable) inductors can be used then the "Exact" values are preferred. These values were obtained from the program *SVC Filter Designer* which is available with the online content.

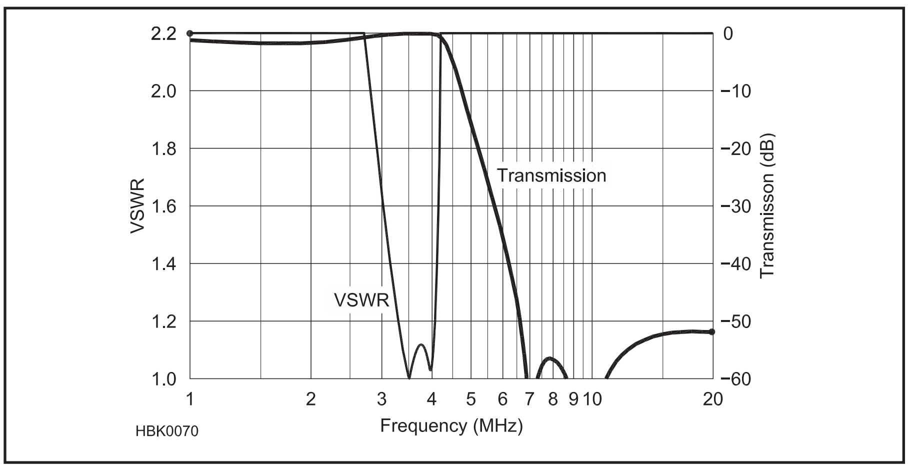

Figure 10.77 — Responses of the filter shown in Figure 10.71. Note the low values of SWR from 3.5 to 4 MHz. At the same time the harmonics are attenuated to meet regulations. Responses for the other amateur bands are very similar except for the frequency scaling.

Table 10.5
Values for the Optimized Harmonic Filters

Band (meters)	*C1 (pF)*	*L2, 5% (µH)*	*L2, Exact (µH)*	*C2 (pF)*	*C3 (pF)*	*L4, 5% (µH)*	*L4, Exact (µH)*	*C4 (pF)*	*C5 (pF)*
160	2400	3.0	2.88	360	4700	2.4	2.46	820	2200
80	1300	1.5	1.437	180	2400	1.3	1.29	390	1100
60	910	1.0	1.029	120	1600	0.91	0.8897	270	750
40	680	0.75	0.7834	91	1300	0.62	0.6305	220	560
30	470	0.56	0.5626	68	910	0.47	0.4652	160	430
20	330	0.39	0.3805	47	620	0.33	0.3163	110	300
17	270	0.30	0.3063	36	510	0.27	0.2617	82	240
15	220	0.27	0.2615	30	430	0.22	0.2245	68	200
12	200	0.24	0.241	27	390	0.20	0.2042	62	180
10	180	0.20	0.2063	24	330	0.18	0.1721	56	150
6	91	0.11	0.108	13	180	0.091	0.0911	30	82

10.11.2 Combline Filters for 50 – 432 MHz

(This project description is condensed from the paper “Combline Filters for VHF and UHF” by Paul Wade, W1GHZ, at the 41st Eastern VHF Conference in 2015. The full paper is available in the online content available with this book.)

RF pollution is rampant at good portable locations on mountaintops and other high places — anywhere accessible is populated with cell phone towers, TV and FM broadcast stations, two-way radio and pager transmitters, and even amateur repeaters. Most of these are high power, producing signals strong enough to seriously overload the VHF and UHF transceivers we use for contest operation or microwave liaison. The advent of broadband MMIC preamps acerbates the problem. The problem often manifests itself as a very high noise level.

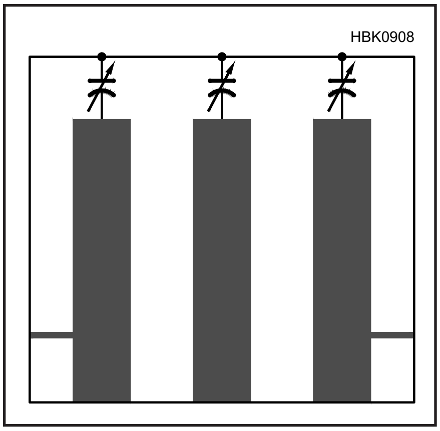

Figure 10.78 — Sketch of combline filter in stripline.

A combline filter uses parallel transmission line resonators less than a quarter-wave long, loaded by capacitance at the open end. This allows tuning over a range of frequencies by varying the capacitance. Typical electrical length of the resonators is between 30 and 60 electrical degrees; a quarter-wavelength is 90 degrees. These filters use stripline construction with tapped input and output coupling, as sketched in **Figure 10.78**, and the design procedure is given in the full paper.

A good, sharp filter must be mechanically robust to stay on frequency, especially for rover work. For low loss, high Q is important, requiring wide striplines with good contact to ground at the bottom, the high-current point. The filter uses all-aluminum construction to prevent dissimilar metal corrosion. All connections are made with #4 tinned solder lugs and stainless-steel hardware, metals that are least likely to interact with aluminum. An inexpensive aluminum enclosure of 220 × 145 × 60 mm was used for the 144 MHz filter in **Figure 10.79**. Dimensions for all of the filters are included in **Table 10.6**. Enclosure model numbers can be searched online to find sources for the boxes. A close match is sufficient.

The full paper includes construction guidelines, including selecting parts, metalworking, and assembly details. Tuning instructions are given and require the use of a sweep generator or a vector network analyzer can be used. The performance curves in **Figure 10.80** are indicative of completed filter performance. If that equipment is not available, the filters can be optimized at one frequency and used as narrowband filters.

10.11.3 Broadcast-Band Rejection Filters

Inadequate front-end selectivity or poorly performing RF amplifier and mixer stages often result in unwanted cross-talk and overloading from adjacent commercial or amateur stations. Two passive receive-only filters are described here — a high-pass, multi-section filter and a simple series wave trap.

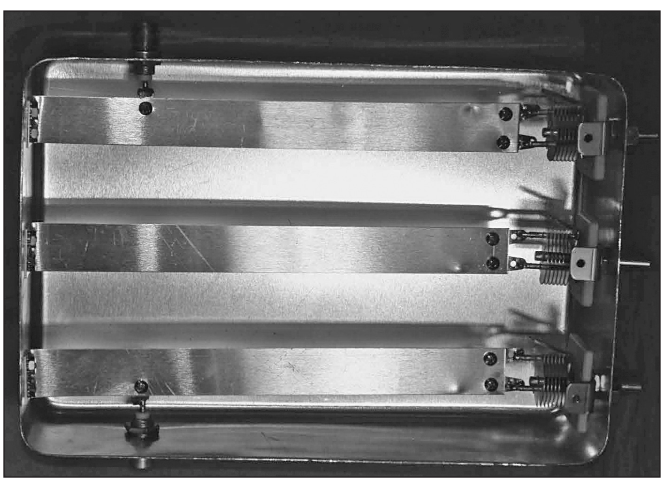

Figure 10.79 — Combline filter for 144 MHz.

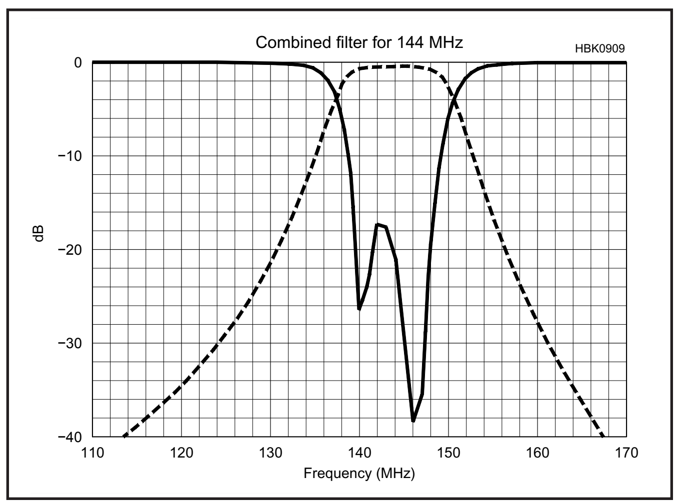

Figure 10.80 — Performance of 144 MHz combline filter. The solid line is Return Loss and the dashed line is Insertion Loss.

Table 10.6
Dimensions for Combline Filter in Stripline

Band (MHz)	*Bandwidth (MHz)*	*BoxModel*	*Length (mm)*		*Width (mm)*		*Depth (mm)*		*Strip*	
width (mm)	*Strip spacing (mm)*	*Strip c to c (mm)*	*Tap point (mm)*		*Capacitor (pF)*					
144	2.5	AC-406	9 in	7 in	2 in	33	44	77	22	24
222	8	U3879 mid	202	129	54	34	40	74	30	15
432	11	1590-BB	115	90	30	16	25	41	16	5
432	13	U3879 sm	176	99	43	29	35	64	15	5
50	3	AC-1418	8 in	10 in	2.5 in	30	40	70	90	150

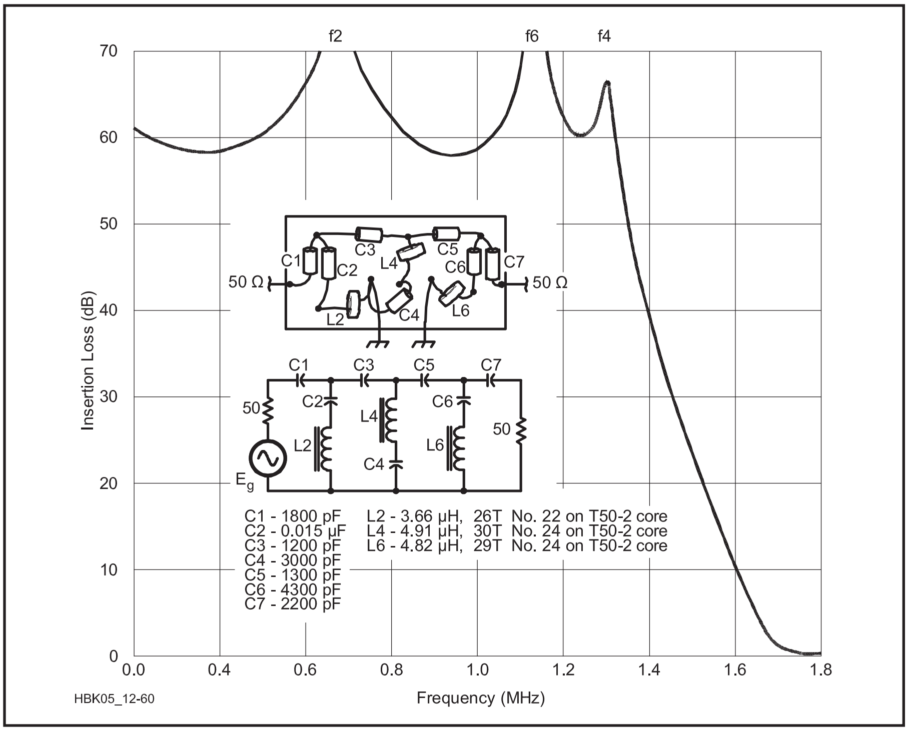

Figure 10.81 — Schematic, layout and response curve of the broadcast band rejection filter.

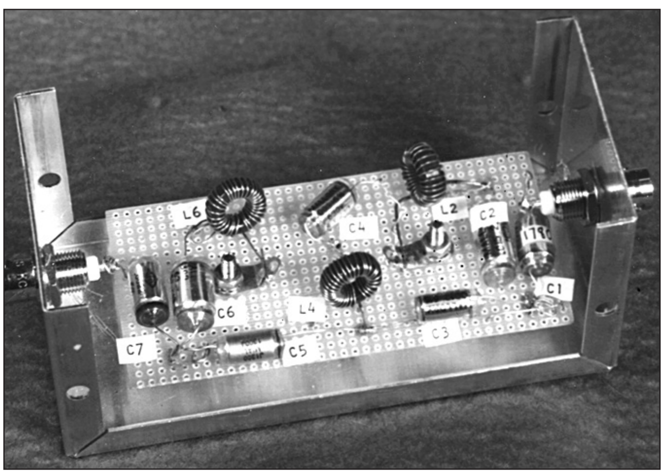

Figure 10.82 — The filter fits easily in a 2 × 2 × 5 inch enclosure. The version in the photo was built on a piece of perfboard.

BROADCAST-BAND REJECTION HIGH-PASS FILTER

The filter shown is inserted between the antenna and receiver. It attenuates the out-of-band signals from broadcast stations but passes signals of interest (1.8 to 30 MHz) with little or no attenuation.

The high signal strength of local broadcast stations requires that the stop-band attenuation of the high-pass filter also be high. This filter provides about 60 dB of stop-band attenuation with less than 1 dB of attenuation above 1.8 MHz. The filter input and output ports match 50 Ω with a maximum SWR of 1.353:1 (reflection coefficient = 0.15). A 10-element filter yields adequate stop-band attenuation and a reasonable rate of attenuation rise. The design uses only standard-value capacitors.

The filter parts layout, schematic diagram, response curve and component values are shown in **Figure 10.81**. The standard capacitor values listed are within 2.8% of the design values. If the attenuation peaks (f2, f4 and f6) do not fall at 0.677, 1.293 and 1.111 MHz, tune the series-resonant circuits by slightly squeezing or separating the inductor windings.

Construction of the filter is shown in **Figure 10.82**. Use polypropylene film-type capacitors. These capacitors are available through Digi-Key and other suppliers. The powdered-iron T50-2 toroidal cores are available through Amidon, Palomar Engineers and others.

For a 3.4 MHz cutoff frequency, divide the L and C values by 2. (This effectively doubles the frequency-label values in Figure 10.81.) For the 80 meter version, L2 through L6 should be 20 to 25 turns each, wound on T50-6 cores. The actual turns required may vary one or two from the calculated values. Parallel-connect capacitors as needed to achieve the nonstandard capacitor values required for this filter.

The measured filter performance is shown in Figure 10.81. The stop-band attenuation is more than 58 dB. The measured cutoff frequency (less than 1 dB attenuation) is under 1.8 MHz. The measured passband loss is less than 0.8 dB from 1.8 to 10 MHz. Between 10 and 100 MHz, the insertion loss of the filter gradually increases to 2 dB. Input impedance was measured between 1.7 and 4.2 MHz. Over the range tested, the input impedance of the filter remained within the 37- to 67.7-Ω input-impedance window (equivalent to a maximum SWR of 1.353:1).

WAVE TRAP FOR BROADCAST STATIONS

Nearby medium-wave broadcast stations can sometimes cause interference to HF receivers over a broad range of frequencies. A wave trap can catch the unwanted frequencies and keep them out of your receiver.

The way the circuit works is quite simple. Referring to **Figure 10.83**, you can see that it consists essentially of only two components, a coil L1 and a variable capacitor C1. This series-tuned circuit is connected in parallel with the antenna circuit of the receiver. The characteristic of a series-tuned circuit is that the coil and capacitor have a very low impedance (resistance) to frequencies very close to the frequency to which the circuit is tuned. All other frequencies are almost unaffected. If the circuit is tuned to 1530 kHz, for example, the signals from a broadcast station on that frequency will flow through the filter to ground, rather than go on into the receiver. All other frequencies will pass straight into the receiver. In this way, any interference caused in the receiver by the station on 1530 kHz is significantly reduced.

This is a series-tuned circuit that is adjustable from about 540 kHz to 1600 kHz. It is built into a metal box, **Figure 10.84**, to shield it from other unwanted signals and is connected as shown in Figure 10.83. To make the inductor, first make a *former* by winding two layers of paper on the ferrite rod. Fix this in place with black electrical tape. Next, lay one end of the wire for the coil on top of the former, leaving about an inch of wire protruding beyond the end of the ferrite rod. Use several turns of electrical tape to secure the wire to the former. Now, wind the coil along the former, making sure the turns are in a single layer and close together. Leave an

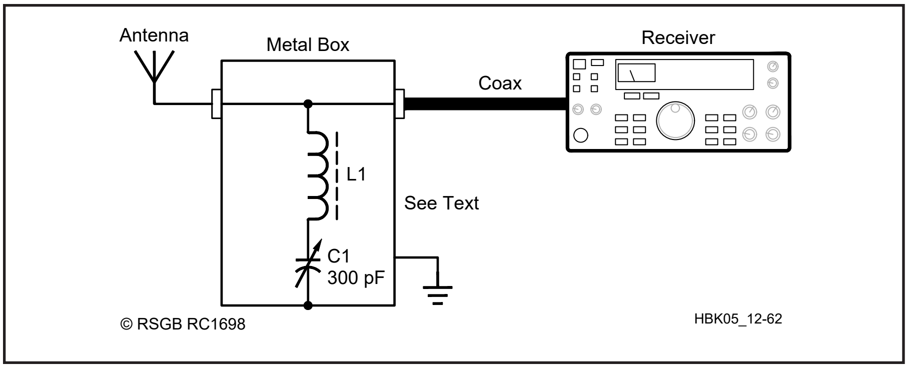

Figure 10.83 — The wave trap consists of a series tuned circuit, which "shunts" unwanted signals on its resonant frequency to ground.

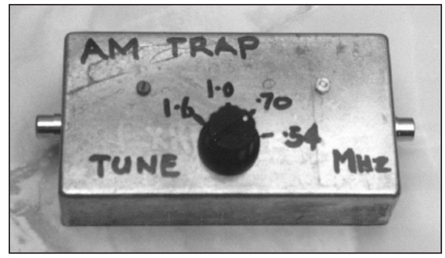

Figure 10.84 — The wave trap can be roughly calibrated to indicate the frequency to which it is tuned.

inch or so of wire free at the end of the coil. Once again, use a couple of turns of electrical tape to secure the wire to the former. Finally, remove half an inch of enamel from each end of the wire.

Alternatively, if you have an old AM transistor radio, a suitable coil can usually be recovered already wound on a ferrite rod. Ignore any small coupling coils. Drill the box to take the components, then fit them in and solder together as shown in **Figure 10.85**. Make sure the lid of the box is fixed securely in place, or the wave trap's performance will be adversely affected by pick-up on the components.

Connect the wave trap between the antenna and the receiver, then tune C1 until the interference from the offending broad-cast station is a minimum. You may not be able to eliminate interference completely, but this handy little device should reduce it enough to listen to the amateur bands. Let's say you live near an AM transmitter on 1530 kHz, and the signals break through on your 1.8 MHz receiver. By tuning the trap to 1530 kHz, the problem is greatly reduced. If you have problems from more than one broadcast station, the problem needs a more complex solution.

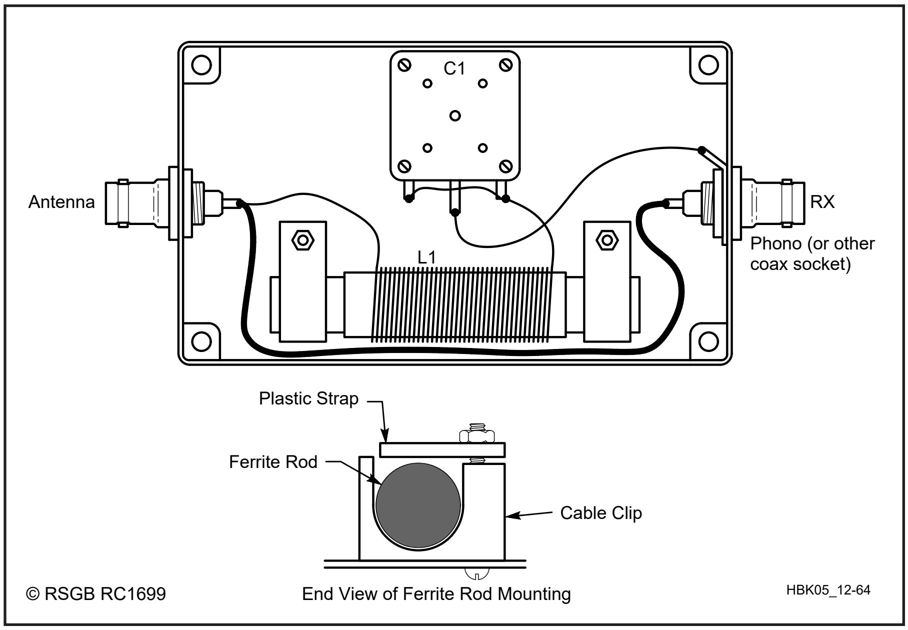

Figure 10.85 — Wiring of the wave trap. The ferrite rod is held in place with cable clips.
C1 — 300 pF polyvaricon variable.
L1 — 80 turns of 30 SWG enameled wire, wound on a ferrite rod.
Associated items: Case (die-cast box), knobs to suit, connectors to suit, nuts and bolts, plastic cable clips.

10.11.4 High-Performance, Low-Cost 1.8 to 54 MHz Low-Pass Filter

The low-pass filter shown in **Figure 10.86** offers low insertion loss, mechanical simplicity, easy construction and operation on all amateur bands from 160 through 6 meters. Originally built as an accessory filter for a 1500 W 6 meter amplifier, the filter easily handles legal limit power. No complicated test equipment is necessary for alignment. It was originally described by Bill Jones, K8CU, in November 2002 *QST*. The complete original *QST* article for this project is included with the online content. The article supplies complete assembly and alignment drawings.

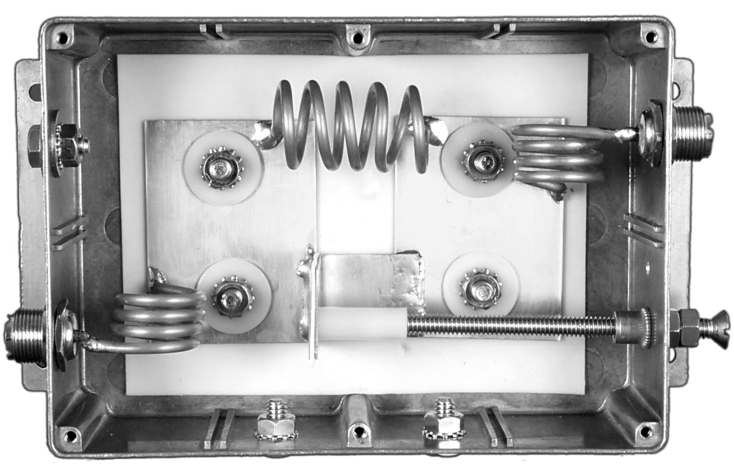

Figure 10.86 — The 1.8 to 54 MHz low-pass filter is housed in a die-cast box.

Although primarily intended for coverage of the 6 meter band, this filter has low insertion

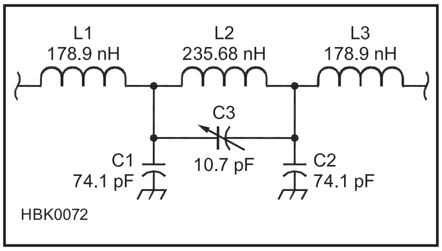

Figure 10.87 — The low-pass filter schematic. See construction details in the article with the online content.

C1, C2 — 74.1 pF. 2 × 2.65 inch brass plate sandwiched with 0.03125 inch thick Teflon sheet. The metal enclosure is the remaining grounded terminal of this capacitor.

C3 — Homemade variable using brass and Teflon.

L1, L3 — 178.9 nH. Wind with ⅛ inch OD soft copper tubing, 3.5 turns, 0.75 inch diameter form, 0.625 inch long, ¼ inch lead length for soldering to brass plate. The length of the other lead to RF connector as required.

L2 — 235.68 nH. Wind with ⅛ inch OD soft copper tubing, 5 turns, 0.75 inch diameter form, 1.75 inches long. Leave ¼ inch lead length for soldering.

loss and presents excellent SWR characteristics for all HF bands. Although harmonic attenuation at low VHF frequencies near TV channels 2, 3 and 4 does not compare to filters designed only for HF operation, the use of this filter on HF is a bonus to 6 meter operators who also use the HF bands. Six meter operators may easily tune this filter for low insertion loss and SWR in any favorite band segment, including the higher frequency FM portion of the band.

ELECTRICAL DESIGN

The software tool used to design this low-pass filter is *Elsie* by Jim Tonne, W4ENE, which is available with the online content. The *Elsie* format data file for this filter, DC54.lct, may be downloaded for your own evaluation from the author's website at **www.realhamradio.com**.

Figure 10.87 is a schematic diagram of the filter. The use of low self-inductance capacitors with Teflon dielectric easily allows legal limit high power operation and aids in the ultimate stop band attenuation of this filter. Capacitors with essentially zero lead length will not introduce significant series inductance that upsets filter operation. This filter also uses a trap that greatly attenuates second harmonic frequencies of the 6 meter band. The parts list for the filter is given in **Table 10.7**.

PERFORMANCE DISCUSSION

Assuming the 6 meter SWR is set to a low value for a favorite part of the band, the worst case calculated forward filter loss is about 0.18 dB. The forward loss is better in the HF bands, with a calculated loss of only 0.05 dB from 1.8 through 30 MHz. The filter cutoff frequency is about 56 MHz, and the filter response drops sharply above this.

Figure 10.88 shows the calculated filter response from 1 to 1000 MHz. The impressive notch near 365 MHz is because of these inherent stray capacitances across each of the coils. Slight variations in each coil will make slightly different tuned traps. This will introduce a stagger-tuned effect that results in a broader notch.

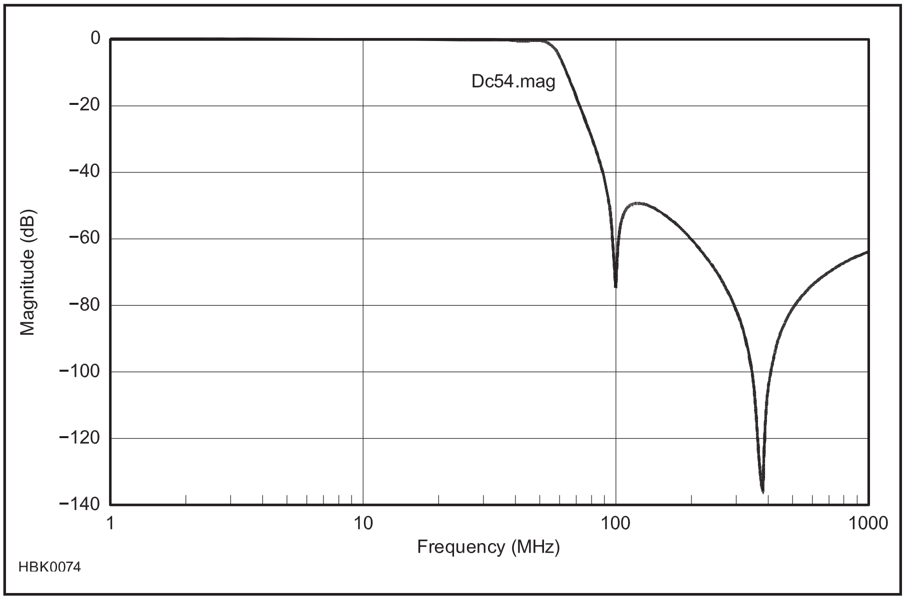

Figure 10.88 — Modeled filter response from 1 to 1000 MHz.

Table 10.7
Low-Pass Filter Parts List

Qty	*Description*
1	Miniature brass strip, 1 × 12 in., 0.032 in. thick (C3)
1	Miniature brass strip, 2 × 12 in., 0.064 in. thick (C1, C2)
5 ft	⅛ inch diameter soft copper tubing
4	¼ × 20 × ½ in. long hex head bolt
4	Plastic spacer or washer, 0.5 in. OD, 0.25 in. ID, 0.0625 in. thick
6	¼ × 20 hex nut with integral tooth lock washer
1	¼ × 20 × 4 in. long bolt
1	¼ × 20 threaded nut insert, PEM nut, or "Nutsert"
1	1 × 0.375 in. diameter nylon spacer. ID smaller than 0.25 in. (used for C3 plunger).
4	Nylon spacer, 0.875 in. OD, 0.25 to 0.34 in. ID, approx. 0.065 in. or greater thickness (used to attach brass capacitor plates).

Aluminum diecast enclosure is available from Jameco Electronics (www.jameco.com) part no. 11973. The box dimensions are 7.5 × 4.3 × 2.4 in. The 0.03125 in. thick Teflon sheet is available from McMaster-Carr Supply Co (www.mcmaster.com), item #8545K21 is available as a 12 × 12 in. sheet.

10.11.5 Band-Pass Filter for 145 MHz

The following project is based on a design from the RSGB *Radio Communication Handbook*, *11th edition.* This filter is intended to reduce harmonics and other out-of-band spurious emissions when transmitting and suppress strong out-of-band incoming signals which could overload the receiver. The filter's schematic and response curve are shown in **Figure 10.89**. This filter design is not suitable for use as a repeater duplexer for in-band signals sharing a common antenna — a high-Q cavity resonator is required.

Direct inductive coupling from the proximity of the coils is used between the first two and the last two resonant circuits. C_C performs "top coupling" between the center two sections where a shield prevents stray

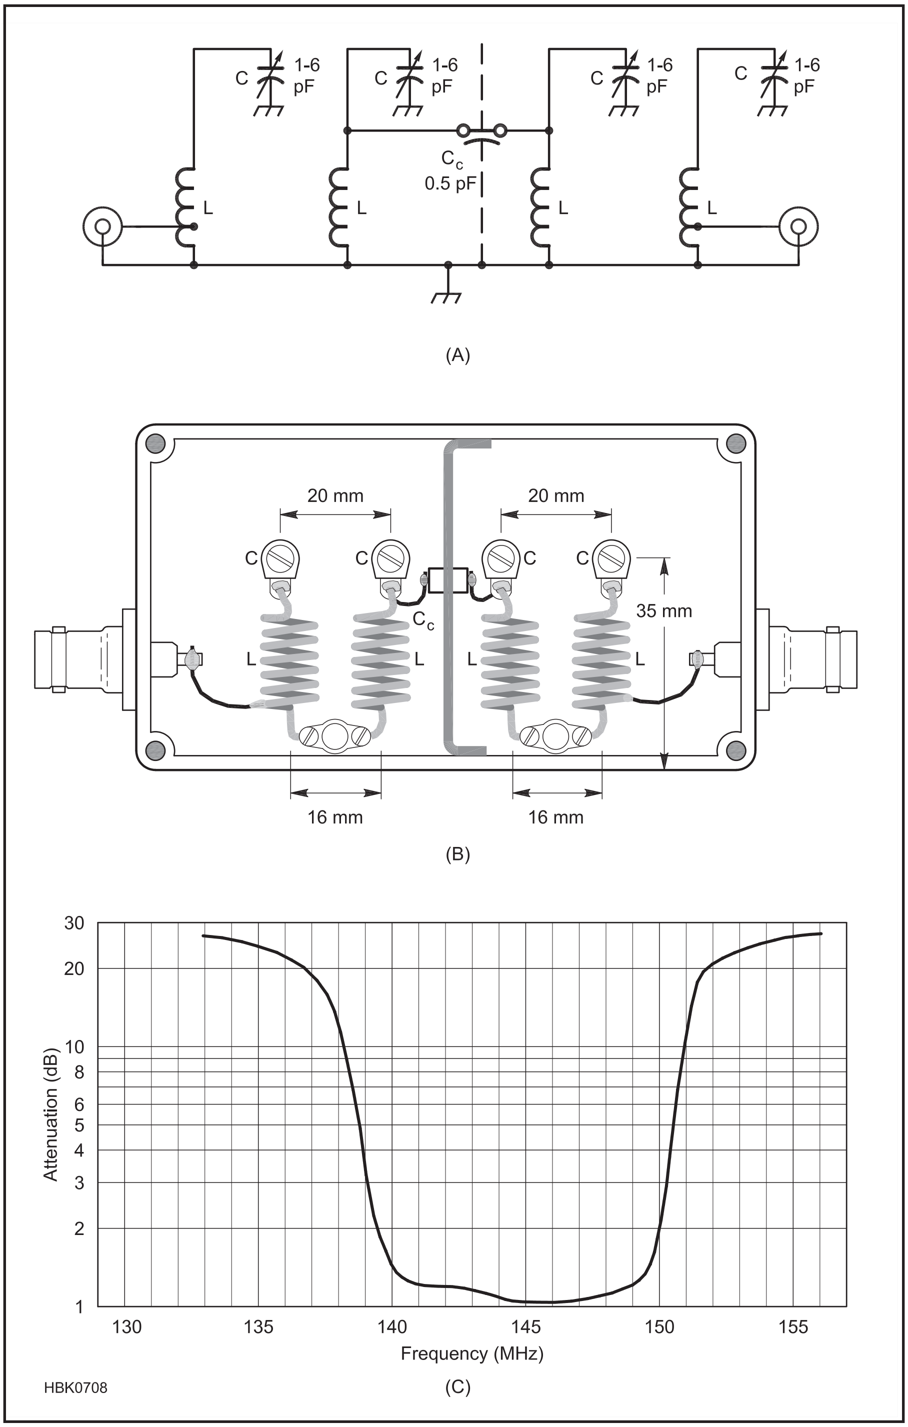

Figure 10.89 — A four-section band-pass filter for 145 MHz for attenuating strong out-of-band signals and reducing harmonics or other out-of-band spurious emissions.

coupling. The input and output connections are tapped down on their respective coils to transform the 50 Ω source and load into the proper impedances for terminating the filter

At VHF, self-supporting coils and mica, ceramic or air-dielectric trimmer capacitors give adequate results for most applications. Piston ceramic trimmers can be used for receiving or for low-power signals of a few watts. At higher power, use an air-variable capacitor. The filter is assembled in a 4 × 2½ × 1½ inch die-cast aluminum box so there is room for either piston or air-variable capacitors.

The band-pass filter is made from four parallel resonant circuits formed by the inductance of the coils and their associated parasitic inter-turn capacitance described in the **RF Techniques** chapter. A 1-6 pF series capacitor (C in Figure 10.89) connected to ground adjusts the resonant frequency of each coil. A ceramic capacitor should be used for C_C which has a value of 0.5 pF.

The dimensions of the coils are important to control the self-resonant frequency. Each coil is constructed from 6½ turns of #17 AWG solid, bare wire (1.15 mm dia), ⅜ inch (9.5 mm) in diameter. Each turn is spaced 1 wire diameter apart. The original design used British #18 SWG wire which is slightly thicker (1.22 mm) so coils made with #17 AWG wire will have a slightly higher resonant frequency. The taps for the coils are made 1 turn from the grounded end of the coil as shown in the figure.

Adjustment will be required after assembly to tune out the stray capacitances and inductances. A sweep generator and oscilloscope (see the **Test Equipment and Measurements** chapter) provide the most practical adjustment method. A variable oscillator with frequency counter and a voltmeter with RF probe, plus a good deal of patience, can also do the job.

10.12 References and Bibliography

References

"A Filter Primer," Maxim Integrated Tutorial 733, 2008. **https://pdfserv.maximintegrated.com/en/an/AN733.pdf**

Bottom, V., *Introduction to Quartz Crystal Unit Design* (Van Nostrand Reinhold Company, 1982).

Butterworth, S., "On the Theory of Filter Amplifiers," *Experimental Wireless and Wireless Engineer*, Oct. 1930, pp. 536 – 541.

Chirlian, P., "Laplace Transforms," *Basic Network Theory* (McGraw Hill, 1969).

Crawford, J. KØZR, "High-Power HF Band-Pass Filter Design," *QEX*, Mar. 2018, pp. 3 – 10.

Cutsogeorge, G., W2VJN, "High-Power Harmonic Filters" *NCJ*, Nov. 1998, pp. 14 – 15.

Darlington, S., "Synthesis of Reactance 4-Poles Which Produce Prescribed Insertion Loss Characteristics," *Journal of Mathematics and Physics*, Sep. 1939, pp. 257 – 353.

Dishal, M., "Top Coupled Band-Pass Filters," *IT&T Handbook*, 4th edition (American Book, Inc, 1956).

Donadio, M., "CIC Filter Introduction" **dspguru.com/dsp/tutorials/cic-filter-introduction**.

Fisk, J., W1DTY, "Helical-Resonator Design Techniques," *QST*, June 1976, pp. 11 – 14.

Geffe, P., *Simplified Modern Filter Design* (John F. Rider, 1963).

Gordon-Smith, D., G3UUR, "Extended Bandwidth Crystal Ladder Filters With Almost Symmetrical Responses," *QEX*, Jul./Aug. 2011, pp. 36 – 44.

"Group Delay Explanations and Applications," Applied Radio Labs, DN004, **www.radio-labs.com/wp/index.php/design-file/**.

Hayward, W., Campbell, R., and Larkin, R., *Experimental Methods in RF Design*, (ARRL, 2003).

Hogenauer, E., "An Economical Class of Digital Filters For Decimation and Interpolation," *IEEE Transactions on Acoustics, Speech and Signal Processing*, Vol. 29, Apr. 1981 pp.155 – 162.

Lyons, R., "Understanding Cascaded Integrator-Comb Filters," *EE Times*, Mar. 31, 2005.

Lyons, R., *Understanding Digital Signal Processing* 3rd Edition (Pearson, 2011).

Mack, R., W5IFS, "SDR: Simplified" Jan./Feb., Mar./Apr., May/Jun., Jul./Aug., and Nov./Dec. *QEX* 2009.

Metcalf, W., "Graphs Speed Interdigitated Filter Design," *Microwaves*, Feb. 1967.

Munson, D., "Fourier Waveform Analysis," *Reference Data for Engineers* 7th Edition, Jordan, E., Ed., (Howard W. Sams & Co., 1985).

Rhea, R., *HF Filter Design and Computer Simulation* (McGraw-Hill, 1995).

Saal, R. "Cauer Elliptic Filters" *The Design of Filters Using the Catalogue of Normalized Low-Pass Filters*, (AEH Telefunken, 1968).

Shepherd, F., "Modern-Network-Theory Design," *Reference Data for Engineers*, 7th Edition, Jordan, E., Ed., (Howard W. Sams & Co., 1985) pp. 9-5 – 9-11.

Silver, W., NØAX, "Hands-On Radio," *QST*, Experiment 50 Mar. 2007, Experiment 51 Apr. 2007, Experiment 87 Apr. 2010, Experiment 88 May 2010. **www.arrl.org/hands-on-radio**

Steder, H., DJ6EV, and Hardcastle, J., G3JIR, "Crystal Ladder Filters for All," *QEX*, Nov./Dec. 2009, pp. 14 – 18.

Weinberg, L., "Network Design by use of Modern Synthesis Techniques and Tables," *Proceedings of the National Electronics Conference*, 1956, Vol. 12.

White, D., *A Handbook on Electrical Filters: (Synthesis, Design and Applications)* (Don White Consultants, Inc., 1980).

Williams, A., *Electronic Filter Design Handbook* (McGraw-Hill, 1981).

Zobel, O., "Theory and Design of Electric Wave Filters," *Bell System Technical Journal*, Jan. 1923.

Zumbahlen, H., "Allpass Filters," Analog Devices, Mini-tutorial MT-202.

Zumbahlen, H., "The Bessel Response," Analog Devices, Mini-tutorial MT-204.

Zumbahlen, H., "Biquadratic (Biquad) Filters," Analog Devices, Mini-tutorial MT-205.

Zumbahlen, H., "The Butterworth Response," Analog Devices, Mini-tutorial MT-224.

Zumbahlen, H., "The Chebyshev Response," Analog Devices, Mini-tutorial MT-206.

Zumbahlen, H., "F_0 and Q in Filters," Analog Devices, Mini-tutorial MT-210.

Zumbahlen, H., "Low-Pass to Band-Pass Filter Transformation," Analog Devices, Mini-tutorial MT-215.

Zumbahlen, H., "Low-Pass to Band-Stop (Notch) Filter Transformation," Analog Devices, Mini-tutorial MT-216.

Zumbahlen, H., "Low-Pass to High-Pass Filter Transformation," Analog Devices, Mini-tutorial MT-217.

Zumbahlen, H., "Multiple Feedback Filters," Analog Devices, Mini-tutorial MT-220.

Zumbahlen, H., "Passive and Active Filtering," Analog Devices, AN-281.

Zumbahlen, H., "Sallen-Key Filters," Analog Devices, Mini-tutorial MT-222.

Zumbahlen, H., "State Variable Filters," Analog Devices, Mini-tutorial MT-223.

Zumbahlen, H., "Twin T Notch Filter," Analog Devices, Mini-tutorial MT-225.

Zumbahlen, H., "Using the Analog Devices Active Filter Design Tool," Analog Devices, AN-649, **https://www.analog.com/media/en/technical-documentation/app-notes/an-649.pdf**.

Zverev, A., *Handbook of Filter Synthesis* (John Wiley and Sons, 1967).

Practical Filter Designs

Antoniazzi, P., IW2ACD, and Arecco, M., IK2WAQ, "Easy Microwave Filters Using Waveguides and Cavities," *QEX*, Sep./Oct. 2006, pp. 37 – 42.

Bartlett, B., VK4UW, and Loftus, J., VK4EMM, "Band-Pass Filters for Contesting," *National Contest Journal*, Jan. 2000, pp. 11 – 15.

Cefalo T., WA1SPI, "Diplexers, Some Practical Applications," *Communications Quarterly*, Fall 1997, pp. 19 – 24.

Cope, P., W2GOM/7, "The Twin-T Filter," *QEX*, Jul./Aug. 1998, pp. 45 – 49.

Fisher, R., W2CQH, "Combline V.H.F. Band-pass Filters," *QST*, Dec. 1968, pp. 44 – 45.

Fisher, R., W2CQH, "Interdigital Band-pass Filters for Amateur VHF/UHF Applications," *QST*, Mar. 1968, p. 32.

Gordon-Smith, D., G3UUR, "Seventh-Order Unequal-Ripple Low-pass Filter Design," *QEX*, Nov./Dec. 2006, pp. 31 – 34.

Gordon-Smith, D., G3UUR, "Fifth-Order Unequal-Ripple Low-pass Filter Design," *QEX*, Nov./Dec. 2010, pp. 42 – 47.

Hayward, W., W7ZOI, "Extending the Double-Tuned Circuit to Three Resonators," *QEX*, Mar./Apr. 1998, pp. 41 – 46.

Heemstra, F., KT3J, "A Double-Tuned Active Filter with Interactive Coupling," *QEX*, Mar./Apr. 1999, pp. 25 – 29.

Jansson, D., WD4FAB, and Wetherhold, E., W3NQN, "High-Pass Filters to Combat TVI," *QEX*, Feb. 1987, pp. 7 – 8 and 13.

Lau, Z., W1VT, "A Narrow 80-Meter Band-Pass Filter," *QEX*, Sep./Oct. 1998, p. 57.

Lumachi, R., WB2CQM, "How to Silverplate RF Tank Circuits," *73 Amateur Radio Today*, Dec. 1997, pp. 18 – 23.
Moliere, T., DL7AV, "Band-Reject Filters for Multi-Multi Contest Operations," *CQ Contest*, Feb. 1996, pp. 14 – 22.
Rahe, W., DC8NR, "High Performance Audio Speech Low-Pass and CW Band-Pass Filters in SVL Design," *QEX*, Jul./Aug. 2007, pp. 31 – 39.
Sabin, W., WØIYH, "Diplexer Filters for the HF MOSFET Power Amplifier," *QEX*, July/Aug 1999, pp. 20 – 26.
Sabin, W., WØIYH, "Narrow Band-Pass Filters for HF," *QEX*, Sep./Oct. 2000, pp. 13 – 17.
Sutcliffe, VK3TCT, "A BPSK band-pass filter," *Amateur Radio*, May 2012, pp. 44 – 46.
Tonne, J., WB6BLD, "Harmonic Filters, Improved," *QEX*, Sep./Oct. 1998, pp. 50 – 53.
Wetherhold, E., W3NQN, "Band-Stop Filters for Attenuating High-Level Broadcast-Band Signals," *QEX*, Nov. 1995, pp. 3 – 12.
Wetherhold, E., W3NQN, "Clean Up Your Signals with Band-Pass Filters," *QST*, Part 1, May 1998; Part 2, June 1998.
Wetherhold, E., W3NQN, "CW and SSB Audio Filters Using 88-mH Inductors," *QEX*, Dec. 1988, pp. 3 – 10.
Wetherhold, E. "CW Filter Combines High Performance and Low Cost" *Radio Handbook*, 23rd Edition, Orr, W., Ed., (Howard W. Sams and Co., 1987) p. 13-4.
Wetherhold, E., W3NQN, "A CW Filter for the Radio Amateur Newcomer," *Radio Communication*, Jan. 1985, pp. 26 – 31.
Wetherhold, E., W3NQN, "Modern Design of a CW Filter Using 88 and 44-mH Surplus Inductors," *QST*, Dec 1980, pp. 14 – 19. See also Feedback in Jan. 1981 *QST*, p. 43.
Wetherhold, E., W3NQN, "Receiver Band-Pass Filters Having Maximum Attenuation in Adjacent Bands," *QEX*, Jul./Aug. 1999, pp. 27 – 33.
Wetherhold, E., W3NQN, "Second-Harmonic-Optimized Low-Pass Filters," *QST*, Feb. 1999, pp. 44 – 46.