Using Active Filter Design Tools

Editor's note: Section and figure references in this article are from the 2010-2013 editions of the ARRL Handbook. *This material was originally contributed to the* Handbook *by Dan Tayloe, N7VE.*

Sophisticated active filter circuits are more easily designed using filter-design software. Follow the same general approach described in Chapter 11 of *The ARRL Handbook* to determine the filter's performance requirements and then the filter family. 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 in this document shows how a real analog design is assembled by understanding the performance requirements and then using design software to experiment for a "best" configuration.

This design example makes use of Texas Instrument's "freeware" filter design software, *FilterPro*. This package is extremely useful in designing active RC filters. This package allows filter parameters to be adjusted so as to "tweak" the design close to standard component values. (Go to **www.ti.com** and search for "FilterPro".) 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.

DESIGN EXAMPLE: 750-HZ HIGH-PERFORMANCE DIRECT CONVERSION RECEIVER FILTER

This example illustrates the design of the front-end filter for a high-performance direct-conversion I-Q phasing CW receiver. Along with the design of the filter, there will be some discussion of how much gain to assign to each circuit in a sequence of stages. This is included to illustrate some of the processes by which performance requirements for a circuit are established.

A unity-gain filter with a fixed-gain block in front of it means that the fixed-gain block will have no out-of band signal rejection and thus may be subject to overload due to strong out-of-band signals. Conversely, a unity-gain filter followed by a fixed-gain block could degrade signal sensitivity due to the noise internal to the unity-gain filter section. A filter with distributed gain is better than either of these two situations as the signals can be both filtered and amplified stage by stage.

The input signal for this example will be assumed to be straight from the detector with a 100-Ω output impedance and may be as large as 4 V_{P-P} (+16 dBm). This filter is assumed to operate from 12-V supplies so that up to 8 V_{P-P} can easily be handled by commonly available op amps such as a LM5532. (Unless an op amp is capable of rail-to-rail output, its output voltage can approach no closer than 2 V from either the positive or negative supply rails, thus 12 – 2 – 2 = 8 V of total output swing.)

A typical receiver requires approximately 80 to 90 dB of total gain for adequate headphone output levels. Speaker-level output usually requires an additional 20 dB of gain. If the volume control is placed too early in the gain chain (at the antenna input, for example), the audio stage will always be running at maximum gain and the result will be a high level of unwanted internally generated receiver hiss, noise that is not affected by the volume control. On the other hand, if the volume control is too late in the gain chain (such as at the receiver output), all the signals being received will be amplified by the maximum receiver gain of 80 to 90 dB. This can cause strong signals in the passband to saturate the audio chain, causing unwanted distortion unless the receiver uses AGC to reduce the gain.

A reasonable compromise is to place the volume control roughly halfway along the chain of gain stages. Thus, if 80 dB of total gain is desired, roughly 40 dB would occur before the volume control and 40 dB after. This keeps the gain after the volume control low enough that unwanted receiver hiss will be largely eliminated when the volume control is turned all the way down as long as a low-noise amplifier chain is used. This implies that the first 40 dB of gain will be before the volume control, which in this design will be rolled into the active RC filter.

The design objectives at this point are:

1) To provide 40 dB of gain in the desired 750 Hz passband, but 0 dB of gain at 2 kHz and higher.

2) At every stage in the filter, the signal at 2 kHz should not be allowed to exceed 8 V_{P-P} for a 4-V_{P-P} input. This means gain at 2 kHz should be 20 log (8/4) or 6 dB or less out of each stage.

3) The filter should be designed to minimize the impact on receiver sensitivity. It should not add unnecessary noise that would mask weak signals or be susceptible to overload from strong signals.

The first and second goals are attempts to ensure that no signal out of the detector at 2 kHz or higher will overload the filter section. This enables the creation of a very high performance direct-conversion receiver. The second goal is to allow that receiver to have high sensitivity which in turn has noise implications on the filter design.

Audio Filter Q Implications

When using the *FilterPro* software, Q must be specified for each filter section. In experimenting with the filter types of the Bessel, Butterworth, and Chebyshev and observing the frequency responses, it can be quickly seen that higher Q is associated with sharper filter frequency rolloff. From this observation, the conclusion could be drawn that high Q in an active RC filter is a good thing. This conclusion is not entirely true.

From a receiver design perspective, the goal is to reject undesired signals to the highest degree possible. An ideal 750-Hz low-pass filter would pass all signals at and below 750 Hz while completely eliminating all signals 751 Hz and higher. A high-order active RC filter with high-Q filter sections comes closest to this ideal. However, this sharp frequency rolloff does not come without a price.

The first problem with high-Q filter sections is ringing. **Fig 11.47** was generated using *FilterPro* for a 5th-order Chebyshev filter with 1 dB of ripple, a cutoff frequency of 750 Hz, and 40 dB of gain. This filter provides 20 dB of attenuation at 2 kHz exceeding our design goal of 0 dB. However, notice in particular the sharp peak filter group delay response at the 750 Hz cutoff frequency. This sharp group delay peak is associated with audio ringing.

The effect of ringing in a filter is much the same as ringing a bell. Strike a bell with a hammer, and the bell "rings" at a certain frequency. Likewise, when noisy, static-filled band noise hits a high-Q filter such as the one shown above, this impulse noise tends to produce an audible "ring" sound at the frequency of the delay spike. The effect of the filter ringing is that it actually creates audible interference that interferes with and can mask the desired signal.

It should be noted that simple crystal ladder filters used in many simple superheterodyne or "superhet" receivers have a band-pass characteristic. Thus, there is both a high and a low band-pass edge where group delay peaks occur. That means the typical narrow 300- to 500 Hz-wide CW crystal filter tends to ring badly at both a high (top end of the band-pass response) and a low frequency (bottom end of the band-pass response) at the same time, which makes the ringing audio artifacts twice as bad.

The second problem with high-Q filters is an effect that is not at all obvious. It is simply the fact that our ears do not like them. This is caused by both the high phase and delay varia-

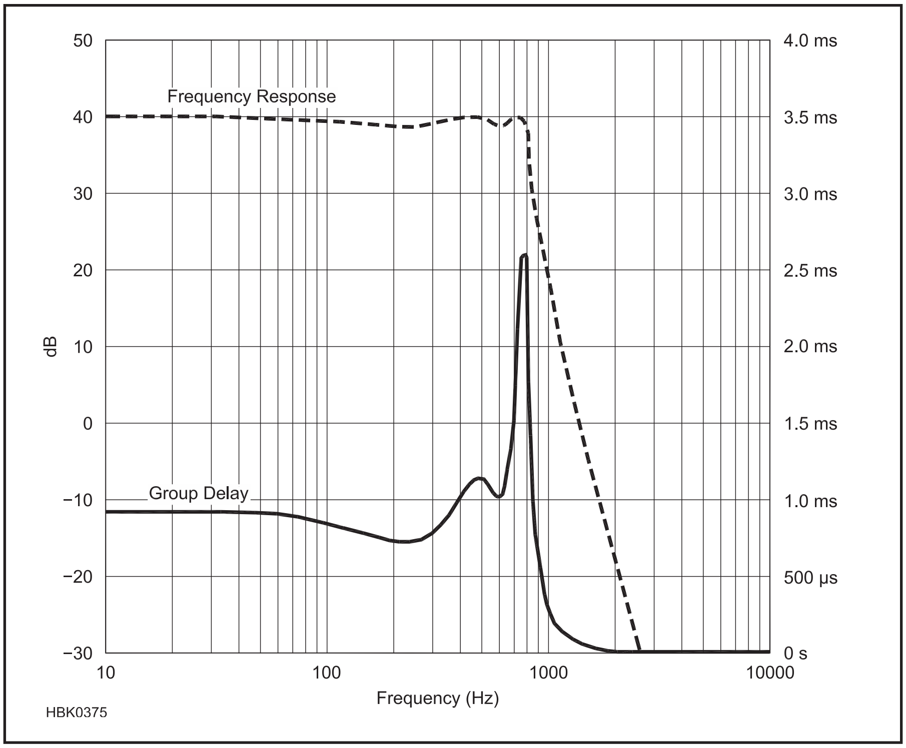

Fig 11.47 — Frequency response and group delay of 5th order Chebyshev, 1 dB passband ripple, 40 dB of gain.

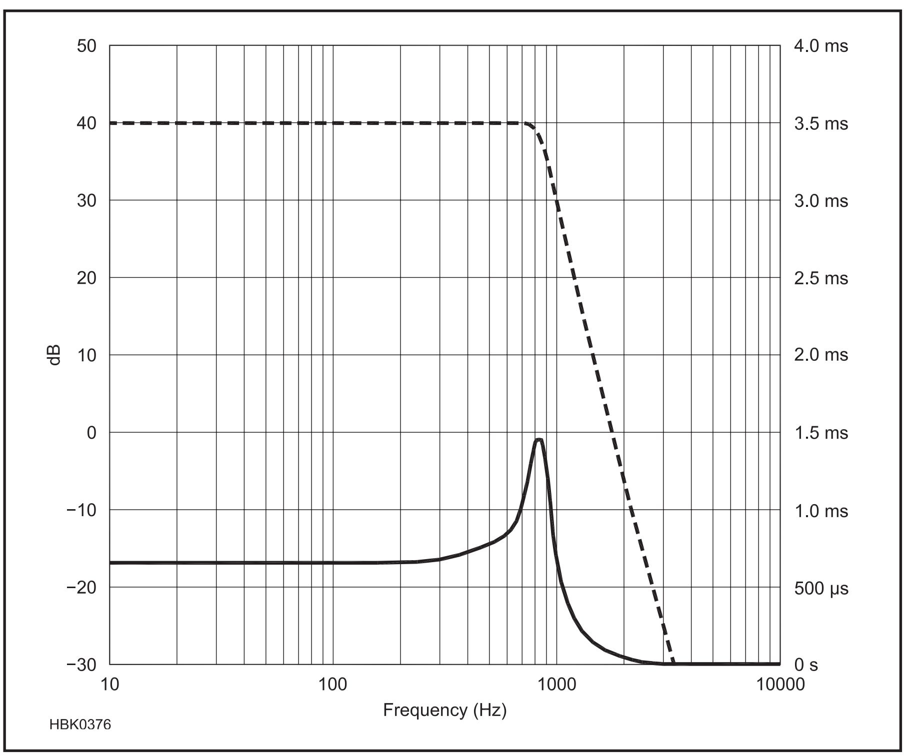

Fig 11.48 — Frequency response (dashed line) and group delay (solid line) of 5th order Chebyshev, 0.06 dB passband ripple.

tions near the edge of the filter. Something in our brain "notices" these variations and objects to them. To our ears, a filter than has lower phase and delay variations "sounds" better, even outside the ringing issue.

In practice, the problems associated with both ringing and phase or delay variations can be reduced by limiting the highest Q filter section to a maximum Q of around 3. When using a design tool like *FilterPro*, this can be done by selecting a filter type of Chebyshev, and then reducing the specified allowed ripple dB value until the Q is reduced to roughly 3. For a 5th-order Chebyshev, this means reducing the allowable ripple from 1 dB to 0.06 dB. The group delay and frequency response of such a filter are shown in **Fig 11.48**.

Notice that the delay peak (the lower line) is now both smaller in amplitude and much broader than that in the previous example, which is exactly what we are after. Notice also that the frequency rolloff (the top curve) is not as sharp either, but it is still 8 dB better than our target of 0 dB at 2 kHz and above.

What does this mean? In active RC filter design, there is a tradeoff between simple and useful vs. more complex and better sounding. A higher-Q active RC filter may require only three filter sections while a lower-Q, better-sounding active RC filter with similar rolloff characteristics may require four filter sections. It is very valuable to realize that a tradeoff is possible and that an active RC filter can be built which is both sharp and sounds good (low phase/delay variations) at the same time.

Noise Implications

Resistors create noise that can mask the small signals we are trying to filter. If this filter is being used at the high-signal end of an audio chain, high-value resistors can be selected without any real harm. Resistance values as high as 1 MΩ can be useful in allowing the selection of small value capacitors which are readily and cheaply available in 2% or 5% tolerance values. However, if this filter is to be used in the front end of a receiver chain, the resistor values need to be much lower. (Receiver noise is also discussed in the **Receivers** chapter.)

A 50-Ω resistor creates about $0.85 \text{ nV} / \sqrt{\text{Hz}}$ of noise. Think of this as the noise generated by a 50-Ω antenna system. This noise voltage varies with the square root of the resistance change. Thus, a 1-MΩ resistor produces $0.85\sqrt{1{,}000{,}000 / 50}$ or $120 \text{ nV} / \sqrt{\text{Hz}}$ of noise. Thus, using a 1-MΩ resistor in the first stage of an active RC filter would reduce the sensitivity of an ideal receiver by 20 log (120/0.85) or 43 dB which is not good for receiver sensitivity.

If each stage of the active RC filter has gain, the effect of that gain is to lessen the impact of the noise contributed by the resistors in each

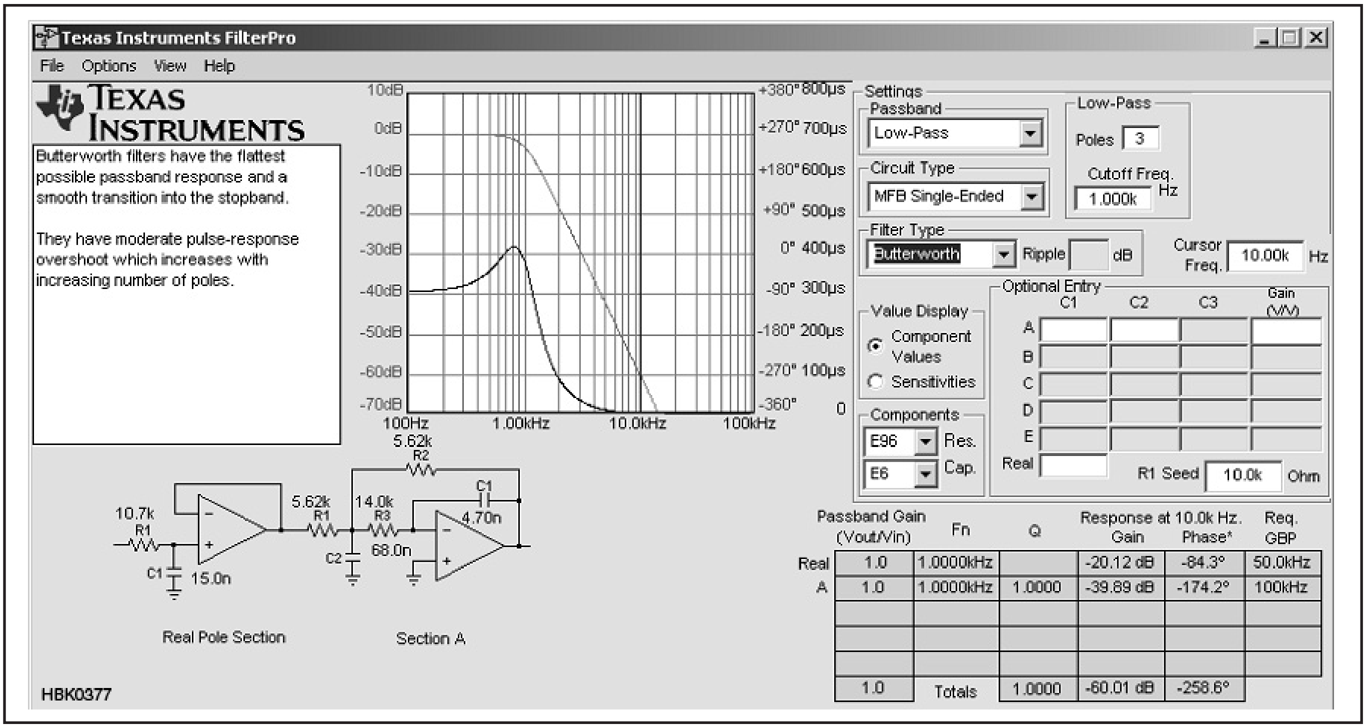

Fig 11.49 — *FilterPro* startup view.

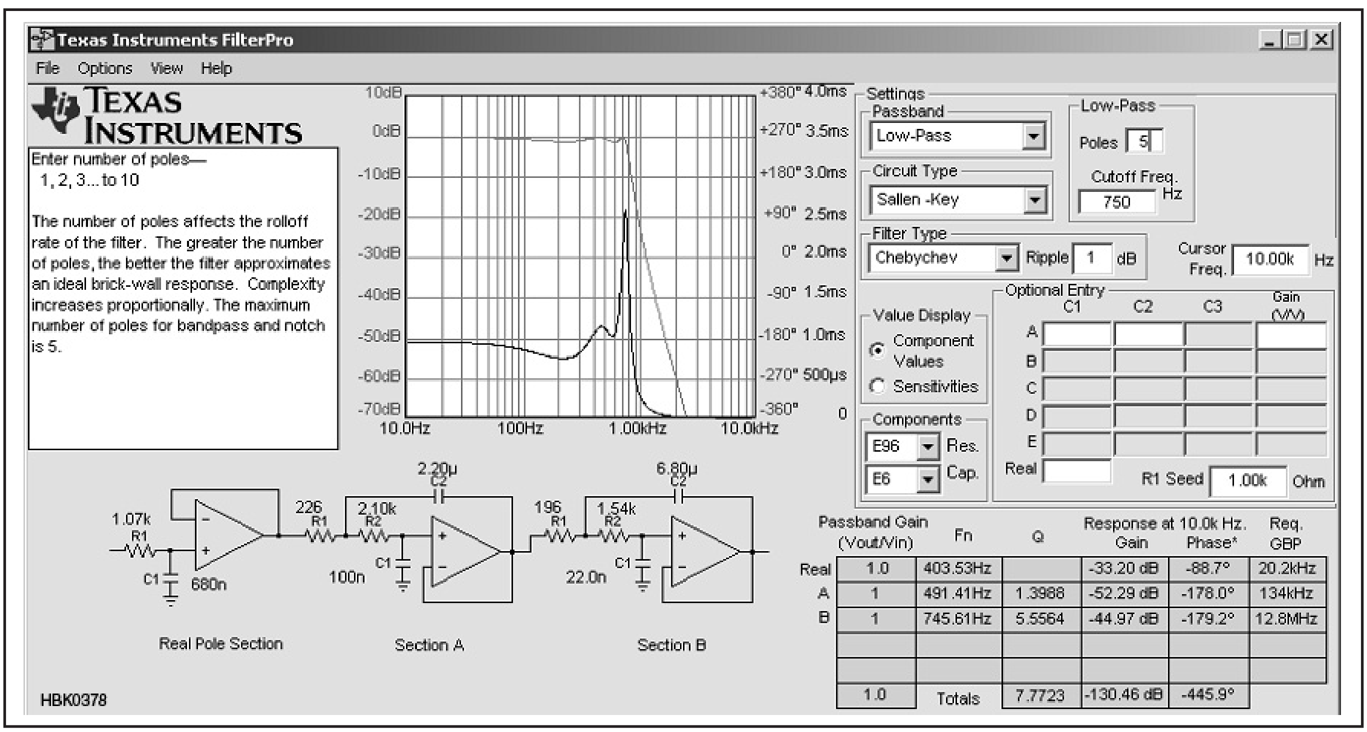

Fig 11.50 — First pass at designing a 5-pole 750-Hz low-pass Chebyshev filter.

succeeding stage. For example, if we want 40 dB (100x voltage gain) in three filter stages, then if the gain is evenly distributed across all the stages, each stage will have a gain of the cube root of 100 or a gain of 4.6 in each of the three stages. In this case, an input signal at $0.85\ \text{nV}/\sqrt{\text{Hz}}$ into the first stage will be 4.6× larger ($3.9\ \text{nV}/\sqrt{\text{Hz}}$) into the second stage and 4.6× larger ($18.3\ \text{nV}/\sqrt{\text{Hz}}$) into the third stage. One goal would be to use resistors than generate at most half the noise voltage of the desired signal. Thus the second stage ($3.9\ \text{nV}/\sqrt{\text{Hz}}$) should have resistors no larger than $3.9 = 0.85\sqrt{x/50}$ or $50 \times (3.9/0.85) \times (3.9/0.85) = 1052\ \Omega$. It should be no problem in the second and third stages to restrict resistor values to the 500-1000-Ω range in order to minimize their noise contribution.

FILTER DESIGN AND COMPONENT VALUE OPTIMIZATION

With an understanding of the gain and component values, we can use *FilterPro* to design our filter. We want a gain of around 4.6× per

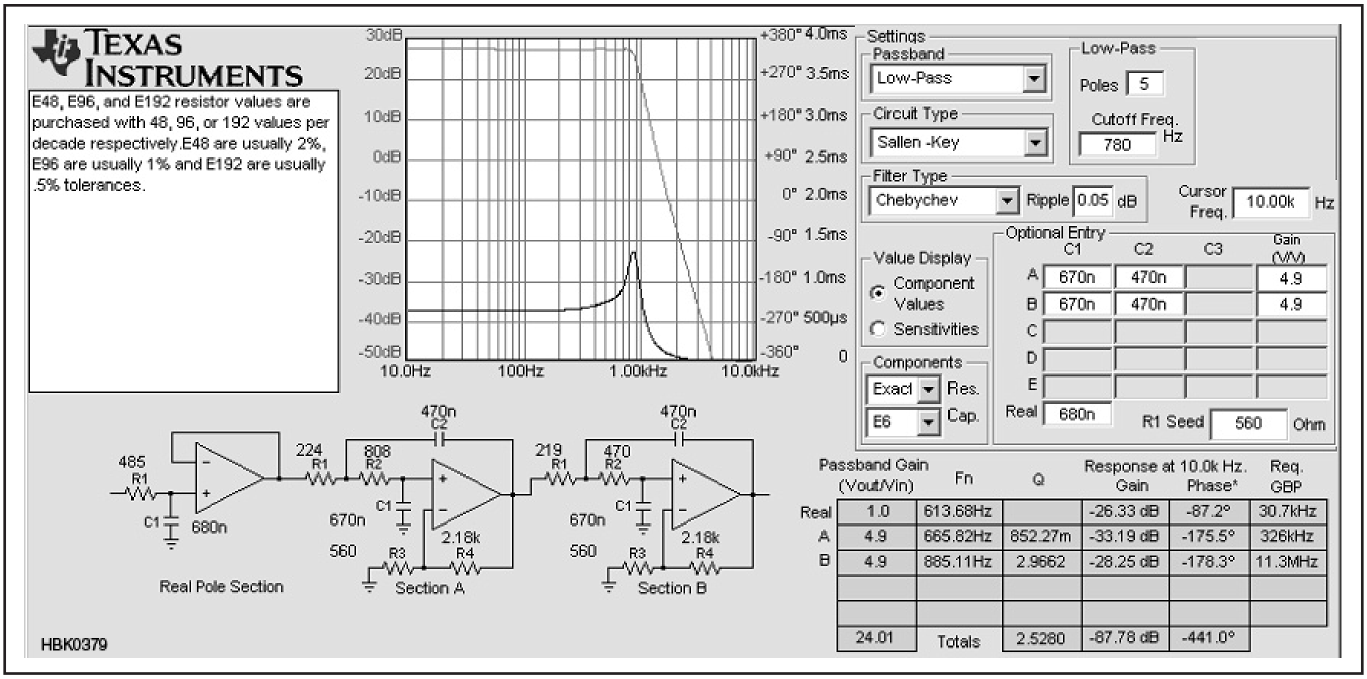

Fig 11.51 — Final result — a 5-pole 780-Hz low-pass Chebyshev filter after component optimization.

stage, resistors in the under-1000 Ω range, a filter Q of 3 or less, and a cutoff frequency of 750 Hz since this is a CW filter. The initial *FilterPro* screen appears as in **Fig 11.49** with a default Passband setting of Low-Pass.

Note the filter frequency response in the graph. Although we have not yet added any gain to any of the stages (we want to add 40 dB or 100×), we can see that the attenuation at 2 kHz is only 20 dB. If we were to add 40 dB of gain, we would still have 20 dB of gain at 2 kHz. Thus we need a sharper filter than a three-pole Butterworth.

We will now change the Filter type to "Chebychev" (one of several spellings), and the Cutoff Frequency to 750 Hz. Under Components, change the resistor setting to "Exact."

In the particular example of a direct-conversion receiver post-detector low-pass filter, it is best to use an odd number of poles so that the odd one-pole section can be readily configured as a receiver post detector preamplifier stage. (One-pole low- and high-pass filters were introduced in Fig 11.45.) Experimenting with *FilterPro*, it can be seen that even numbered pole filters produce more complex stages. Thus, when a three-pole filter is not good enough, the next step up should be a five-pole filter in this particular application.

When using a five-pole Chebyshev filter with a 750 Hz cutoff frequency and an initial value for R1 of 1 kΩ (remember, we want resistances of 1 kΩ or below), *FilterPro* gives the result shown in **Fig 11.50.**

Notice that at 2 kHz, the filter attenuation is now 60 dB. When we create 40 dB of gain in these filter stages, the filter attenuation will still be 20 dB better than needed. However, in the bottom right hand corner, the Q of each section is given, and the highest Q (section B) is 5.55 — higher than the desired maximum Q of 3. The Q of these stages can be adjusted by changing the passband ripple specification, which is set to 1 dB by default for the Chebyshev response. Manually lower the allowable ripple until the Q goes below 3. Experimentally you will find that lowering the passband ripple from 1 dB down to 0.05 dB produces the desired Q of a bit less than 3.

Next set the stages up for the proper gain. You will notice in Fig 11.50 that Circuit Type has been changed to Sallen-Key. This configuration tends to work a bit better as properly biasing the first stage to half the supply voltage will also dc bias all the stages after it and this configuration tends to produce values that are easier to work with. Earlier it was calculated that each stage needed a gain of 4.6× (4.6 × 4.6 × 4.6 = ~100×). This is not a precise value. We want to use component

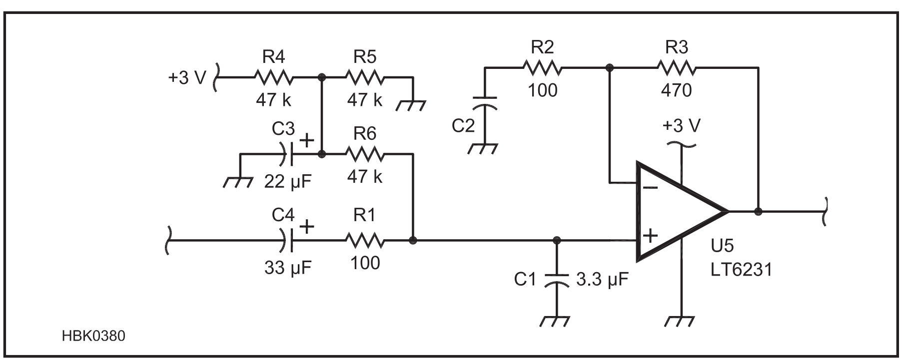

Fig 11.52 — Active filter circuit that implements the "Real Pole" section with biasing and ac decoupling components (see text).

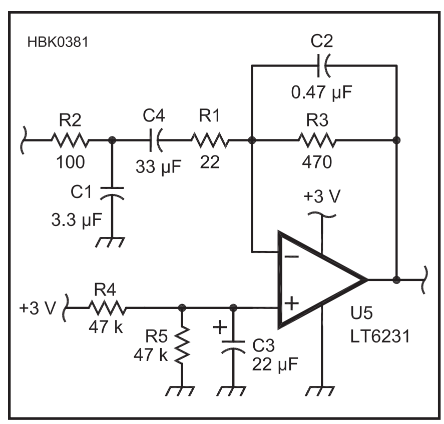

Fig 11.53 — Alternate "Real Pole" circuit with higher performance than Fig 11.52.

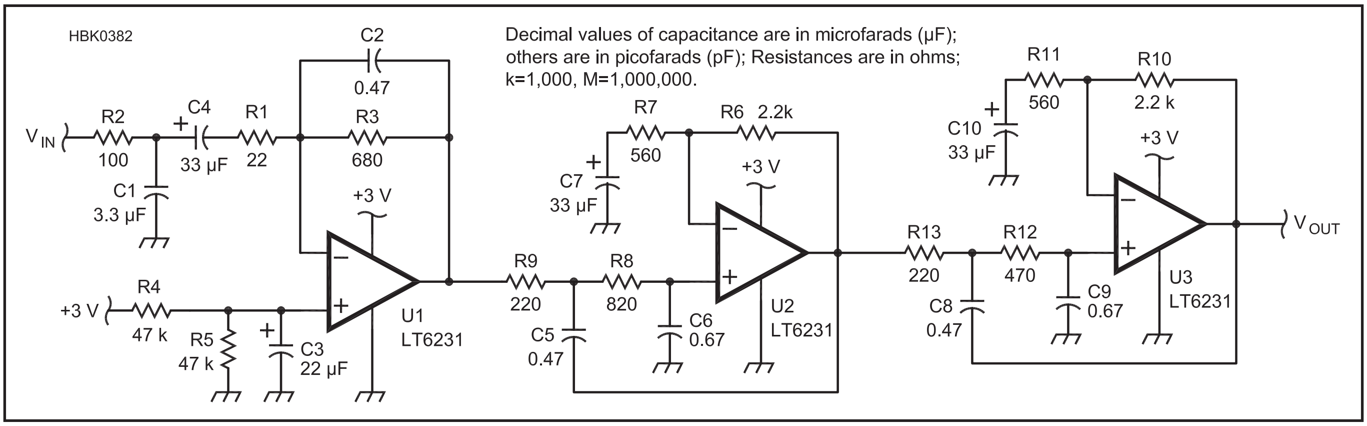

Fig 11.54 — Schematic of the complete filter, including dc biasing components.

values that are close to standard values. Thus we can play with the gain per stage (staying close to 4.6×), and set the "R1 seed" value to obtain values for R3 and R4 close to standard values, then play with C1 and C2 and perhaps adjust the cutoff frequency a bit to get values close to standard values for the other resistors in the two stages. The final result is shown in **Fig 11.51** — a 5-pole, 780-Hz low-pass filter. The software makes the process of "cut and try" much easier than a manual design!

The result comes about from tweaking the cutoff frequency a bit (750 Hz to 780 Hz) and playing with the capacitor values. Capacitor values of 0.67 µF were used instead of 0.68 µF as that value is expensive and hard to find. The 0.67 µF caps are composed of one more-available 0.47 µF capacitor and two very commonly available 0.1 µF capacitors in parallel. This allows the main two sections (section A and B) of the filter to use a total of four 0.47 µF and four 0.1 µF capacitors. In optimizing for component values, section B with higher Q is most sensitive to component value, so section B was optimized to get very close to 220 and 470 Ω (standard values), while the lower-Q section will not be as close when using 220 and 820 Ω.

A note about the gain-setting resistors R3 and R4: The goal was to keep resistors under 1 kΩ in these two stages, but R4 is 2.2 kΩ. As far as noise contributions are concerned, R4 and R3 "look" like they are in parallel to the input of the op amps. Thus, together they look like 1/(1/560 + 1/2180) or 445 Ω, which is indeed much less than 1 kΩ.

This still leaves the one-pole "real pole section" to be configured with the proper gain. There are two approaches that can be taken. The simplest circuit is presented in **Fig 11.52** with R1 = 485 Ω and C1= 680 nF. The input connects to C4 and the output is taken from the output pin of U5.

To match the receiver detector output impedance of 100 Ω, we scale R1 from 485 to 100 Ω. We must also scale C1 up by the same amount (4.85×) to make C1 = 3.3 µF. R1 can actually be eliminated as a separate component, replacing it with the 100-Ω output impedance of the detector. If R1 is eliminated, C1 should be moved to the input side of C4. C1 should be a ceramic type capacitor, not an electrolytic.

If the detector produces its own 1.5-V bias from the 3-V supply (like a Tayloe detector), the bias components R4, R5, R6 and C3 can also be eliminated along with the dc isolation capacitor, C4. U5 is shown as an LT6231 low-noise 3-V op amp. If a less expensive 12 V device is used, such as an LM5532, the bias network will be used to set the bias voltage to 6 V (1/2 the power-supply voltage) and all these parts will be needed.

Another higher-performance variation of the "real pole" section is shown in **Fig 11.53.** Again, R2 can be eliminated if the output impedance of the detector is 100 Ω. In this configuration, C1 and C2 work together to provide rolloff. This configuration provides around 22 dB of attenuation at 10 kHz compared to 12 dB in the first implementation above. The gain of the stage is now R3/(R1+R2) or about 5.5×. Both circuits were simulated using *LTSpice* (free circuit simulation software from Linear Technologies at **www.linear.com/software**), and R3 and C2 were selected (using trial and error) to provide a similar gain peak (13.7 dB vs 13.5 dB) and

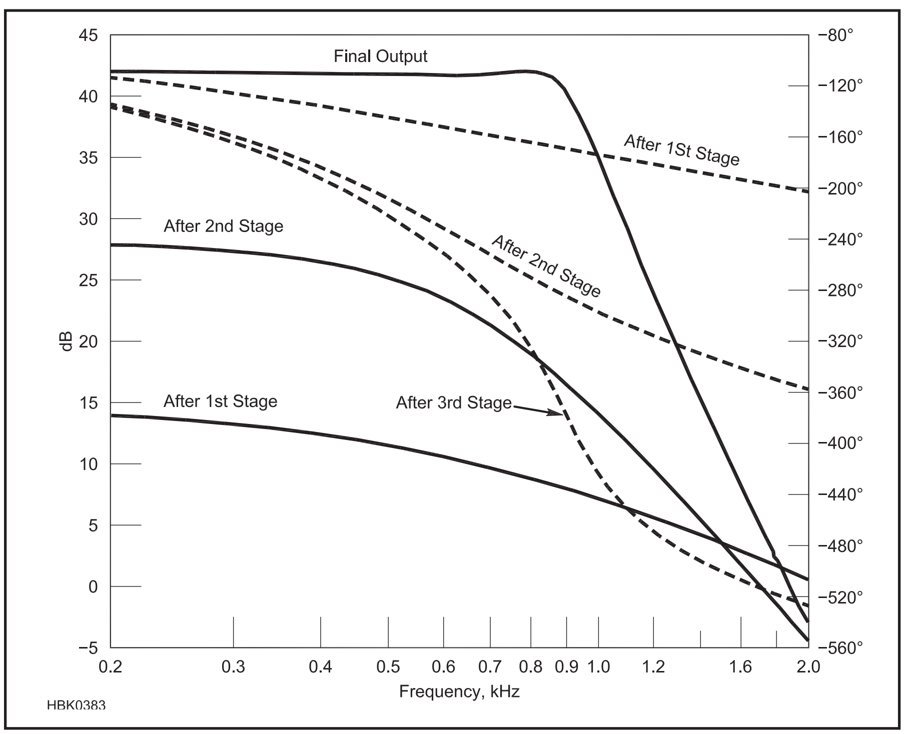

Fig 11.55 — Frequency response of all stages added one at a time. Dashed lines show phase response.

a similar gain at the filter cutoff frequency of 780 Hz. (8.6 dB vs 9.2 dB).

Fig 11.54 is the complete five-pole filter with all necessary biasing and dc isolation, using standard value parts. If a 12-V op amp is used, the 3-V supply voltages shown above will be replaced with 12-V supplies. Remember that the 670 nF capacitors are composed of one 0.47 µF and two common 0.1 µF capacitors in parallel.

The resulting frequency plot (as modeled in *LT Spice)* is shown in **Fig 11.55**. The lowest gain curve is the frequency response of the initial "real pole" section. The gain almost reaches 0 dB by 2 kHz as desired. The next highest gain curve is the frequency response at the output of the second stage (the net response of the first and second stages). Although the total gain is higher, the gain still drops below 0 dB before 2 kHz. The highest gain line shows the total filter response and shows a slight rise near the cutoff frequency. The gain is almost 42 dB (125× voltage gain), close to the desired 40 dB target gain and also drops below 0 dB of gain before 2 kHz.

With this filter placed just after the receiver detector, no signal out of the detector (4 to 5 V_{P-P}) at 2 kHz or higher (just 1.22 kHz above the filter cutoff of 780 Hz) is capable of overloading the front end of the receiver.

Stage Order

In the example above, the stages were ordered from lowest Q (the real section) to the highest Q (section B). This order gives the best protection from inter-stage overload, but is not necessarily the best order for best receiver sensitivity. Fig 11.55 was generated showing the net frequency response adding one stage at a time. However, **Fig 11.56** shows the frequency response of each of the three stages separately.

Notice that the section responses labeled "Real Pole" and "Section A (Q=0.85)" both lose a lot of gain approaching the filter edge at 780 Hz. This loss of gain means that the resistor noise will have more impact than expected, in effect reducing the receiver sensitivity somewhat. Notice that the highest-Q section, labeled "Section B (Q = 3)" actually has a gain peak near the band edge. A peak like this will be present in any filter stage with a Q higher than 1. The higher the Q, the higher and sharper this peak will be.

The "real pole" section has a simpler configuration that makes it easy to use as the first pre-amp stage of the filter, so it comes first. However, moving Section B from the last stage to the second stage will help overcome the stage gain reduction of the first stage and provide for better receiver sensitivity overall as shown in **Fig 11.57**.

When the three stages are ganged in this

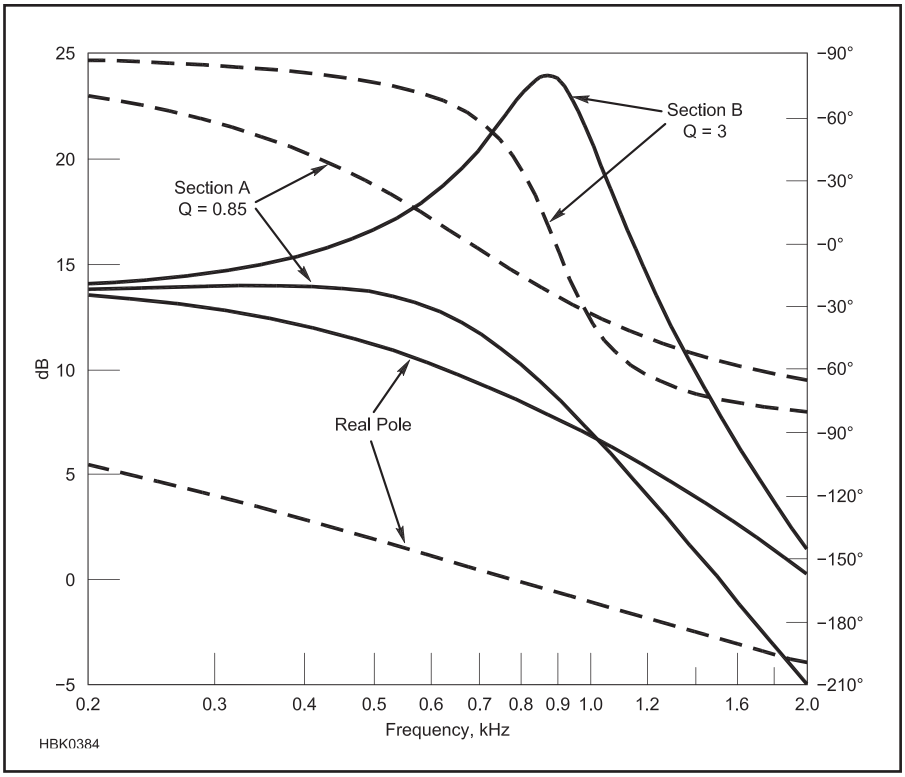

Fig 11.56 — Frequency response of each filter stage shown separately. Dashed lines show phase response.

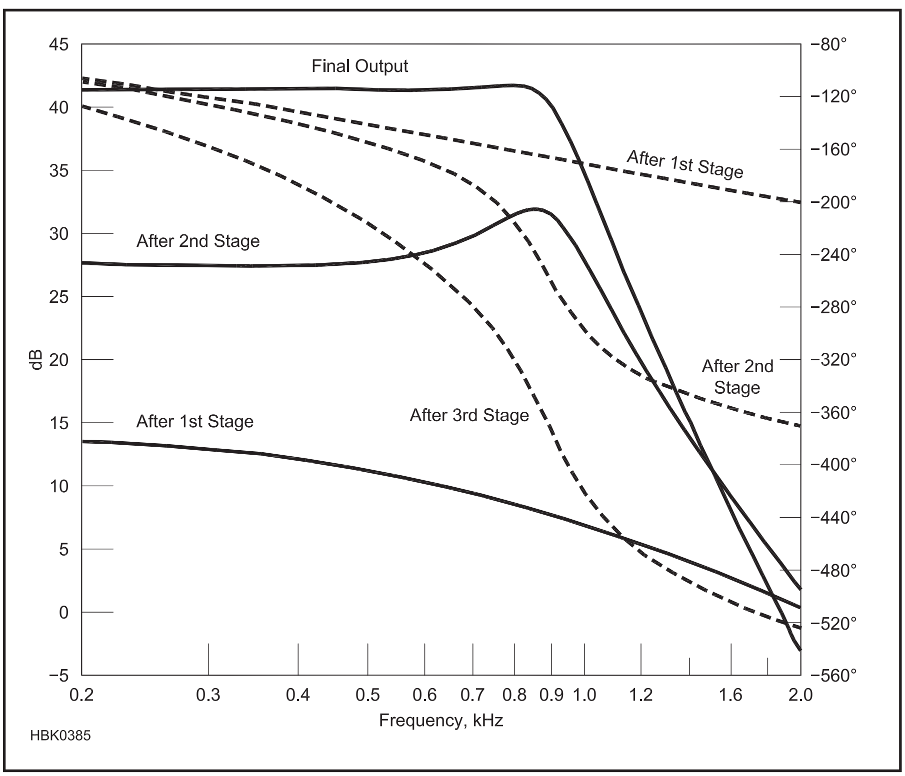

Fig 11.57 — Cumulative frequency response with the high-Q stage in the middle. Dashed lines show phase response.

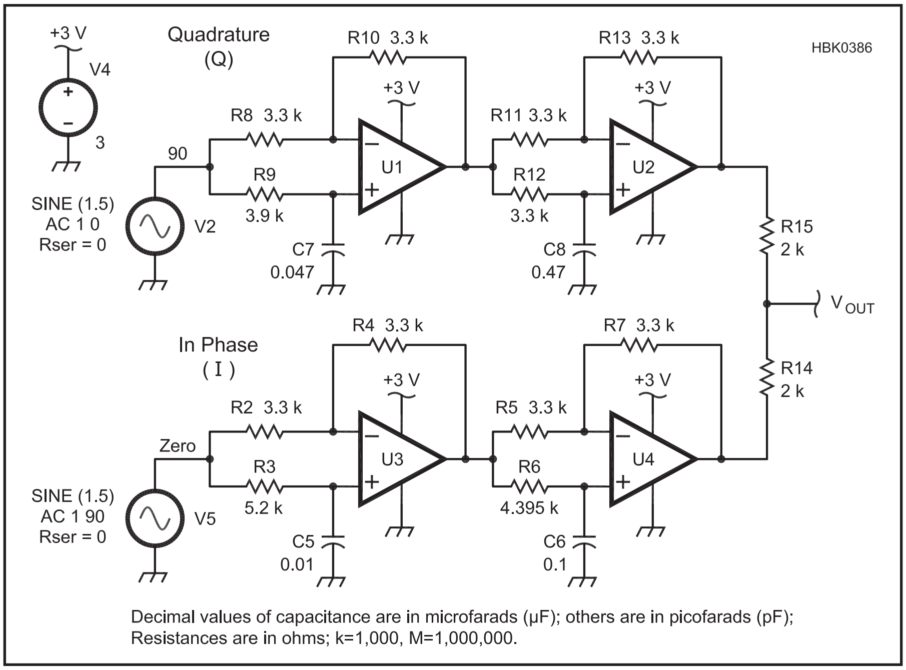

Fig 11.58 — Two all-pass filters create SSB direct-conversion receiver audio from 300 to 1000 Hz.

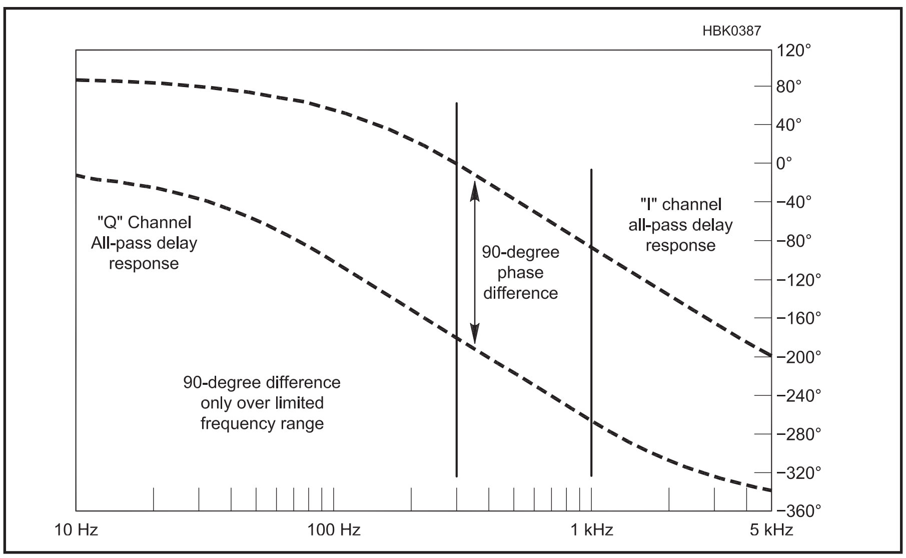

Fig 11.59 — Differential phase delay for two all-pass phasing sections.

order — Real Pole, Section B, Section A — the signal out of the second section is higher near the filter cutoff of 780 Hz, allowing for better low-noise performance. As can be seen in the final output curve, the stage ordering does not change the overall filter frequency response.

A drawback of this configuration could be that the second stage of the filter might now be more susceptible to overload (due to the gain peak) from large signals nearer the edge of the filter than before. This is not really the case in this filter, however. Since the filter was designed with roughly equal gain in each stage, an overload to the second stage is automatically also an overload to the last stage, since it applies additional gain to the second filter section output. However, if this filter had been designed with unity gain per section, the ordering of the stages (and stage overload) could be more of a concern.

In summary, for best ultra-low-noise results (which is only important with very small signals) the stages should be ordered:

- Highest Q sections to lowest Q sections
- If odd order, odd Real Section goes first

For best signal overload protection within the filter bandpass (which is only important if there is little or no gain per section), the sections should be ordered from lowest Q (a real pole section) first (if any) to the highest Q section.

If the filter has many sections and has little or no gain per stage, an optimum balance between good low noise results and good internal passband filter overload protection might be to alternate the sections between the lowest Q and the highest Q sections. For example if a 9-pole filter has a real section, and four other filter sections with Qs of 0.6, 0.9, 1.8, and 5.7, the best section order might be: Real Pole section, followed by Q=5.7, Q=0.6, Q=1.8, and lastly Q=0.9. Simulating and studying the combined frequency response using a program like *LTSpice* is very useful in understanding what is going on.

ALL-PASS ACTIVE FILTERS

A filter that is often seen in phasing-type receivers is the all-pass filter. This filter passes signals of all frequencies without affecting their amplitudes, but creates a controlled phase shift that varies with frequency. This phase shift is used in quadrature direct-conversion receivers for creating single frequency reception by creating a 90° phase difference used to cancel out an unwanted sideband. (See the **Receivers** and the **DSP and Software Radio Design** chapters for more information on direct-conversion receivers.)

The circuit of **Fig 11.58** takes the quadrature I and Q outputs of a direct-conversion quadrature detector (like a Tayloe detector) and adds two all-pass filters designed to create 90° of phase difference between the top and bottom two-stage all-pass sections. U1 and U2 form one phase delay section while U3 and U4 form the other. The 90° difference in phase between these two unity gain all-pass sections is shown in **Fig 11.59.**

In a direct-conversion quadrature receiver, the I and Q outputs are 90° apart from each other and contain the audio from both sidebands. For one of the sideband signals, the I output is 90° ahead of the Q output and for the other sideband signal, Q is 90° ahead of I. Adding in an additional 90° of delay will cause I and Q to both have the same delay on one sideband (–90 + 90 = 0) and 180° of phase difference (90 + 90 = 180) on the other sideband. Since the output is taken via the sum of R14 and R15, a 180° difference will cause the signals from one sideband to cancel, while a 0° difference will allow the signals of the other sideband to add together.

The 90° difference holds well only over a limited range, so suppression of the opposite sideband gets worse at the high and low end of the “sweet spot” where the signals are

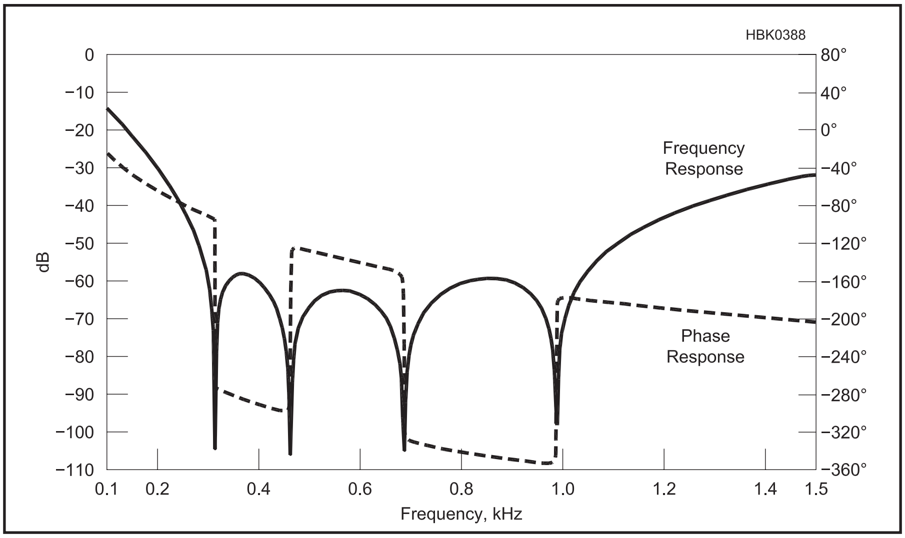

Fig 11.60 — More than 50 dB of opposite sideband suppression is obtained by using the differential phase shift between two all-pass phasing stages.

precisely 90° apart and suppression of the unwanted sideband is best. When adjusted properly, opposite sideband suppression can be excellent over a limited range as shown in Fig **11.60** This was a phasing section designed for CW covering 300 Hz to 1 kHz with well over 50 dB of opposite sideband rejection. A 300-Hz high-pass filter and a 1000-Hz low-pass filter can be used to attenuate the high and low frequency ranges in which sideband suppression drops below 50 dB.

It is difficult to hand-pick components to the 0.1% precision needed to get 60 dB of suppression. Small trimmer resistors can be placed in series with fixed values at R3 and R6 in order to allow the filter to be tuned. The book *Experimental Methods in RF Design* discusses all-pass phasing sections in great detail. The program *QuadNet* for designing and analyzing active quadrative networks is included on the *Handbook* CD-ROM.