L. B. Cebik W4RNL
5105 Holston Hills Road
Knoxville TN 37914

Painless Op-Amp Filter Design

Custom applications can be easy. Just follow this step-by-step guide to a perfect triple op-amp filter.

The triple op-amp audio filter has become a standard, not only in amateur circles but in commercial design as well. Easy to design and nearly foolproof in construction, the various configurations of this filter have found their way into a large percentage of existing ham shacks, either hidden within a transceiver or sitting on the speaker as an audio adjunct. Numerous small companies offer post-receiver audio units using from one to eight filter units.

Even though popular, op-amp filters seem to confuse most ham builders. Despite the low cost of parts, few hams build their own. A simple but effective single filter with a bandpass of between 100 and 200 Hertz would cost about $10.00 for parts, excluding the case and power supply, which together might double the cost. This is a small investment in selectivity, considering what one might learn in the process. Still, there are few takers.

Part of the problem stems from the volume of material that has been written about triple op-amp filters. There are at least three semi-distinct configurations of these filters, but only two different models. However, because designers recast schematic diagrams in different ways, the average ham comes to believe there may be dozens of models. Going even further, different designers choose different circuit values without explaining their choices; the variations seem to grow without limit and without any clear sense. The available books on filter design mire the ham builder in theoretical design math while simultaneously claiming simplicity. There is *some* necessary math to designing a personally-satisfying triple op-amp filter, but it is straightforward hand-calculator stuff.

For the CW buff, most of the existing designs have limitations. Many are fixed-frequency units allowing no tuning to please the ear. The units that permit tuning tend to cover 300 to over 3000 Hz, a fine range for the SSB fan who can use high- and low-pass capabilities built into the filter, but extraneous for CW. A filter that covers a span ranging from 300 to 400 Hz at the bottom to perhaps 1200 Hz at the top would reach two goals. First, the filter would cover the main receiver passband for CW, which runs (depending upon preference) from 400 to 800 Hz wide. Second, the filter would spread its narrower tuning range across the filter frequency dial, permitting the operator to find more easily the desired signal. Unfortunately, most home-brew designers have merely guessed their way into a tuning range.

There is a very direct and easy-to-follow procedure for designing triple op-amp filters in the ham shack. Not only will the procedure ensure a filter that works, but also it will allow the builder to refine the filter's tuning range to his desires. The following notes present a procedure used to design several dozens of different filters for experimental, evaluative, and operational use, and those who have tried the procedure claim they

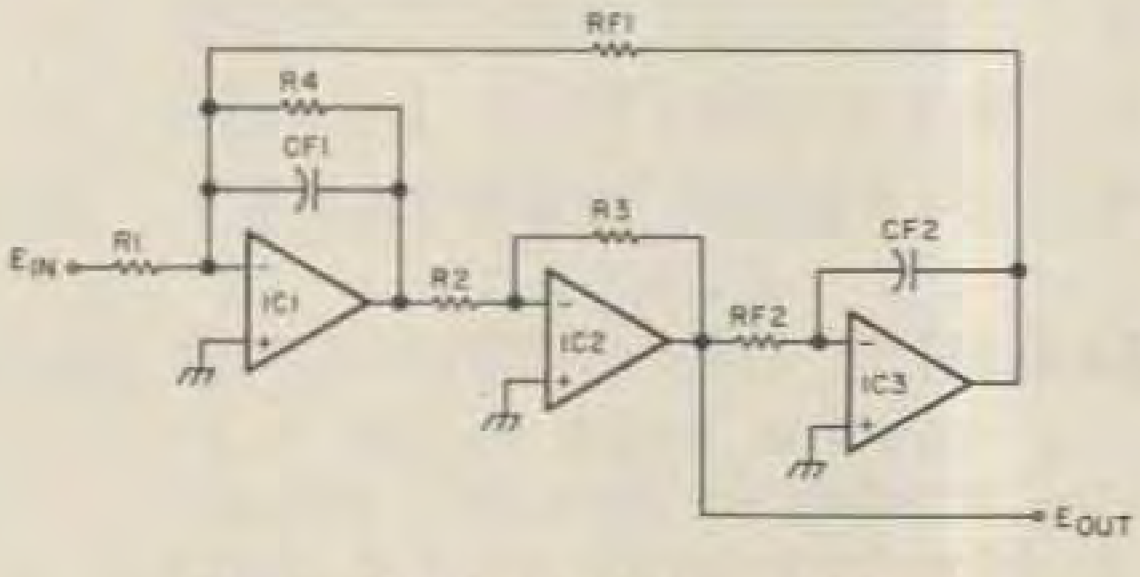

Fig. 1. The basic bi-quad filter.

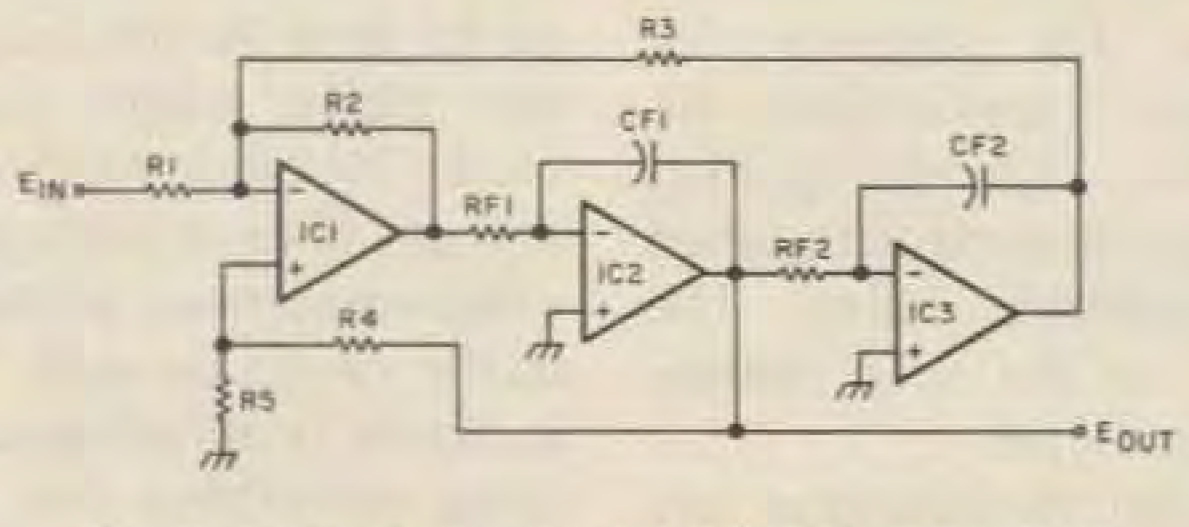

Fig. 2. A basic state-variable filter (—SVF).

have finally made a filter that works and that they like. The procedure even includes steps that show how to let a hand calculator do most of the work.

Some Op-Amp Basics

There are many triple op-amp filter designs but only two fairly distinct types. Unfortunately, the history of op-amp filter terminology has obscured the subject. Originally, the mathematical methods of designing filters gave rise to the name "bi-quad" as a label for all designs. Newer derivations yielded the name "state-variable filter." For some, these names refer only to the design methods; for others, they refer to circuit configurations. At the risk of arousing the wrath of some professional designers, let's follow the latter course.

The bi-quad (or B-Q) appears in Fig. 1. Note that the input op amp is an integrator, as is the third op amp. (Theory aside, an integrator circuit is little more than an op amp whose feedback is provided by a capacitor rather than a resistor.) The middle op amp is an inverter, and we take our bandpass output from this stage. Feedback from the first and third stages is fed back to the first stage input. By controlling the amount of feedback from one of these stages, the first, we control both the gain and the Q or selectivity of the filter. The components marked RF1, CF1, RF2, and CF2 control the frequency of the filter.

Fig. 2 shows the other triple op-amp filter design. The state-variable filter (or −SVF, with the minus sign to be explained very soon) also consists of two integrators, but this time in positions two and three, with a summing amplifier as the input stage. Feedback from the integrators combines with the input signal at the inverting or negative input of the first op amp. We control the gain and Q of the filter by the ratio of resistors R4 and R5, and we set the frequency by the components marked to correspond to those in the B-Q filter. Bandpass output comes from the middle stage, this time an integrator. Unlike the B-Q filter, the −SVF design provides high-pass and low-pass outputs, but at different signal levels than the bandpass output.

The −SVF filter has a near twin which we can call the +SVF. Fig. 3 shows the configuration. The major difference between the SVFs is that this version feeds the input signal to the non-inverting or positive input of the summing op amp. (The reason for the labels +SVF and −SVF should now be clear.) Gain and Q feedback also return to this pin, now being controlled by the ratio of R4 to R1. Although this filter belongs in the SVF family, some of its components require slightly different values from its brother, and the gain vs. Q characteristics will differ. Otherwise, it works perfectly well.

The B-Q and SVF filters have different properties that, for various needs around the shack, give one advantages over the other. First, both SVF filters will have a constant Q and gain throughout their tuning ranges. This means that the bandwidth, when measured in Hertz, will increase as the filter frequency increases. In contrast, the B-Q filter has a constant bandwidth in Hertz, but consequently increases in Q and gain with frequency. For fixed-frequency filters, this phenomenon is meaningless, but for tunable filters, it is important. The constant output of the SVF designs makes follow-up amplification simple. However, every SVF section (i.e., three op-amp filter) requires a dual potentiometer to change RF1 and RF2 together.

The B-Q filter is tunable in the same way but may also be tuned by changing just RF1. Since, like virtually all other filter sections, these filters will ring if the Q is very high, we can cascade two lower Q B-Q sections for a sharper bandpass using only one dual pot. Dual pots are hard enough to find; four-section pots in audio (log) taper are nearly impossible to come by, being either inaccessible or very expensive (which amounts to the same thing for most of us). A newer variety of op amp, the operational transconductance amplifier (OTA), promises to relieve us of these problems, but few practical ham designs using the device have yet to appear.

Notice that there is no clear winner in the contest between the B-Q and the SVF filters. Rather, we must design around their limitations. For example, we can overcome the gain change of the B-Q filter by making the Q resistor, R4, variable, or by following the filter with a limiting amplifier such as the one in Fig. 4. This is the W4MLE variable-compression version of the N6WA Audio Elixir. (See *73* for September, 1979, p. 116, and November, 1982, p. 32.) Until OTAs become more common, there is no way to solve the multiple-pot problem of cascaded SVFs; however, for most work on CW, a single-section, moderate-Q filter requiring just one dual pot will do wonders. A Q of 20 theoretically yields a half-power bandwidth of just 30 Hz at 600-Hz center frequency. Even allowing for low-precision components, we do not need excessively high Q filters to enhance CW. In practice, design Qs in the range of 15 to 20 will yield −6 dB (half-voltage) bandwidths in the 100-to-120-Hz range for a 600-Hz center frequency.

Designing Your Filter

In Fig. 1 through Fig. 3, components having comparable duties have the same designation. For all designs, the frequency-determining components are the same although differently placed. R2 and R3 provide feedback and can be treated as alike in all three cases. In the −SVF design, R1 equals the feedback resistors, while in the +SVF version, it will be half their value. In the B-Q, the input resistor can equal the feedback resistors

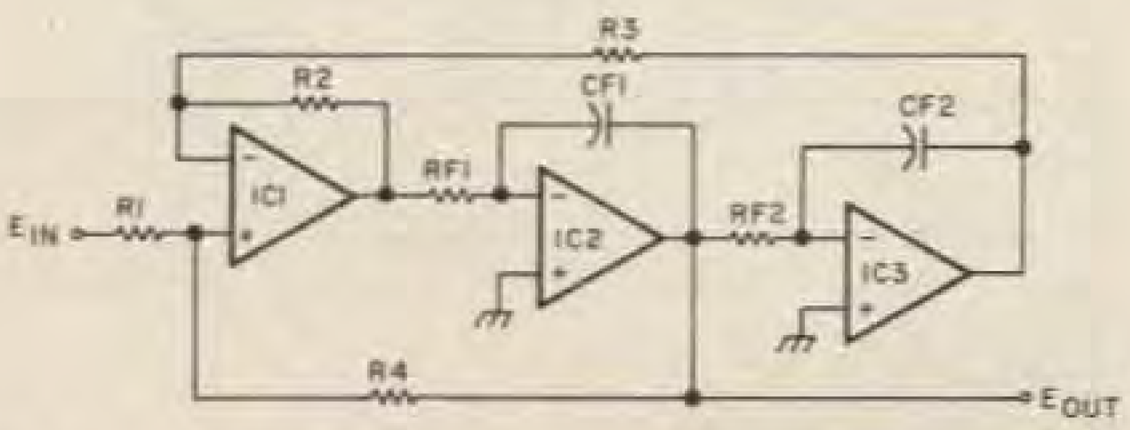

Fig. 3. A basic state-variable filter (+SVF).

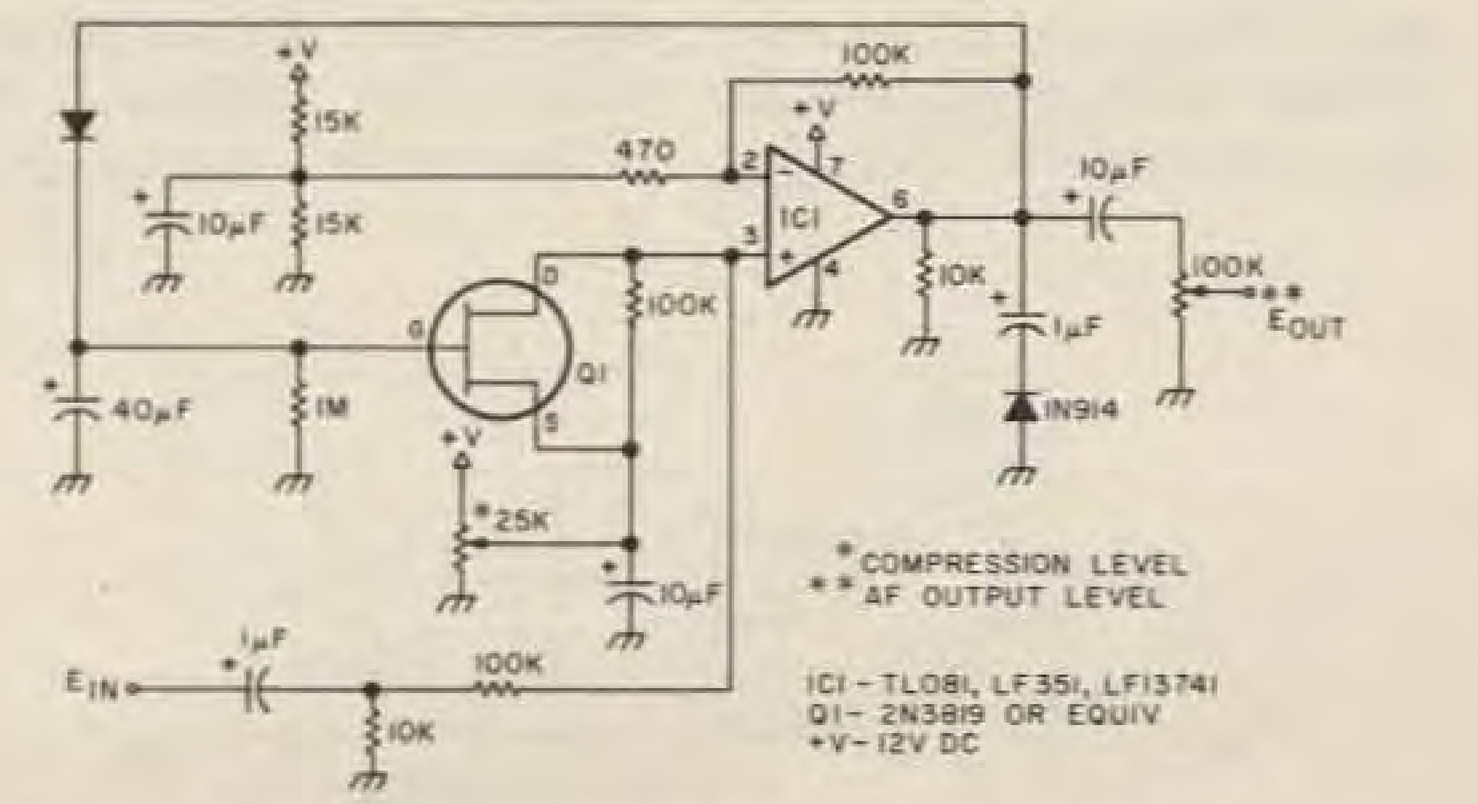

Fig. 4. A limiter/compressor for post-filter amplifying.

Filter Type	State-Variable Inverting Input	State-Variable Non-Inverting Input	Bi-Quad
Schematic	**Fig. 2**	**Fig. 3**	**Fig. 1**
Frequency	$F_c = 1/2\text{pi}R_fC_f$	$F_c = 1/2\text{pi}R_fC_f$	$F_c = 1/2\text{pi}R_fC_f$
Frequency-determining resistors	RF1 = RF2	RF1 = RF2	RF1 = RF2
Frequency-determining capacitors	CF1 = CF2	CF1 = CF2	CF1 = CF2
Bias resistors	R1 = R2 = R3	R2 = R3 = 2R1	R1 = R2 = R3
Q-determining resistors	R4 = R5(3Q − 1)	R4 = R1 (2Q − 1)	R4 = R1Q
Q	Q = (R4 + R5)/3R5	Q = (R4 + R1)/2R1	Q = R4/R1
Gain ($A_o = E_{out}/E_{in}$)	$A_o = Q$	$A_o = 2Q$	$A_o = Q$
Non-inverting input bias resistors	N/A	Fig. 6, Norton amplifier configuration only R6 = R5 = $2R_f$	N/A

Fig. 5. A comparison of filter design relationships.

or vary somewhat from their value according to the needs of the Q relationship. Only in the −SVF design does Q leave the input resistor unaffected, being determined by the relationship between R4 and R5. In the other designs, the input resistor will be a compromise (if needed) between the dictates of Q and the desired situation of having the input resistor correctly related to the feedback resistors.

This discussion may make designing a filter appear difficult. In fact, design is quite easy if done according to a straightforward procedure. Taken step by step, the procedure almost ensures satisfying success. Let's start with some basic relationships, as shown in Fig. 5.

This table reveals where the differences between designs will occur. Calculating R4 will be slightly different for each case. Notice that the +SVF filter has twice the gain of the other designs for a given Q. This may or may not be an advantage. For a filter inserted between the detector and audio amplifier of a receiver, the doubled gain with a low-level input can be useful. For post-receiver use with normal speaker input to the filter, the lower gain of the −SVF and B-Q designs may be more than we need. In all cases, we should have a means of varying the input level.

Aside from these points, design of the three-filter versions will be nearly identical. The first step is to think about the ICs we will use. The LM324 is perhaps standard for both single- and dual-voltage supply applications. Its current requirements are relatively small and it is easy to handle. The TL084 is an FET input version with an identical pinout; its current requirements are even less. The 3900 Norton amplifier also is popular in single-voltage designs, but its biasing is different. Fig. 6 shows the basic configuration of the +SVF design with Norton biasing. Notice the additional formula that sets the values of the bias resistors to the noninverting positive op-amp inputs. Otherwise, our work will be the same as for regular op amps.

Much of the available literature on filters is still written in terms of the relatively high current 741 op amp. Hence, about the highest value shown for feedback resistors is 10k. In fact, 10k should be about the minimum value for R1, R2, and R3. Something approaching 100k is more appropriate, although we will not freeze that value at this point. Instead, we will start by selecting an op amp and the desired frequency range.

This differs from textbook procedures, but for good reasons. First, the ham builder ordinarily has access to components with 5% or 10% tolerances rather than the 1% and .1% tolerances commercial designers prefer. Consequently, absolute peak performance from ham models of op-amp filters is not possible. Very good performance is possible and practical. Since we will aim at good though imperfect performance, we can take a few liberties with absolute precision at some points to gain better precision at points more important to hams.

Second, one of the most evident shortcomings of home-brew filter designs is the fact that tuning controls for frequency and Q rarely cover the most desirable ranges. The techniques for designing filters are easy, but almost never described.

Third, the current crop of op amps available for filter work is very forgiving when we compare the precise operating level to overall filter performance. Hence, we can set our own priorities when establishing a design procedure. In fact, feel free to modify the following procedure to suit personal needs and desires.

While the procedure involves twelve individual steps, they cover only three areas of concern: setting the frequency or tuning range of the filter, ensuring correct feedback, and setting the selectivity and gain of the filter. With a few reservations noted in the procedure steps, these are almost independent design operations. To make the procedure more thoroughly clear, let's step through it, working an example as we go along.

Twelve-Step Filter Design

Step 1. Select an op amp.

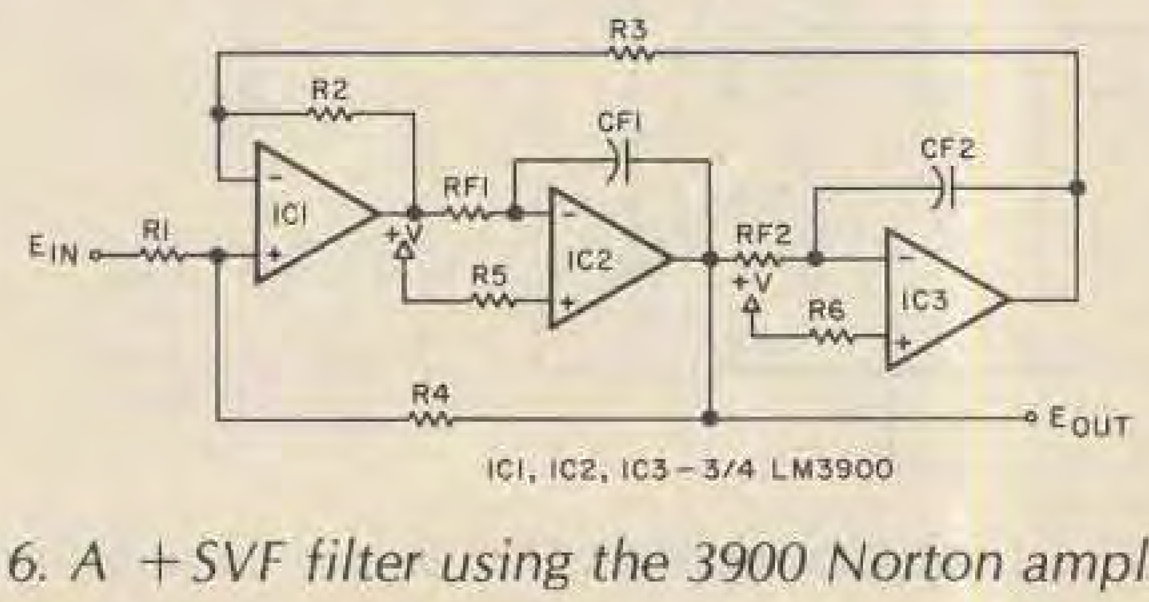

Fig. 6. A +SVF filter using the 3900 Norton amplifier.

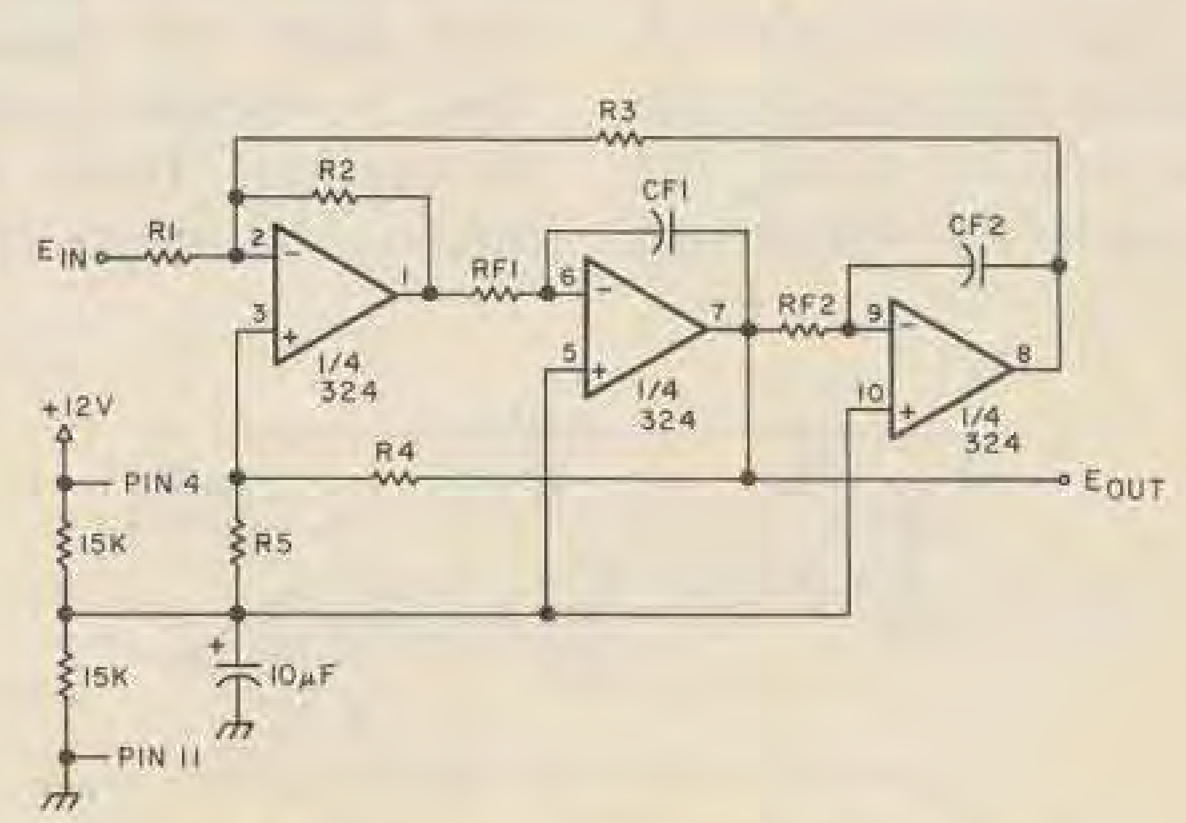

Fig. 7. Setting up the 324 for filter design.

In this case, let's use the reliable LM324.

Step 2. Select a circuit. We can start with the −SVF of Fig. 2 and later see what happens when we convert the design to the other circuits.

Step 3. Choose a power supply. In this example, we have chosen a single 12-volt source. This forces us to provide a voltage divider to feed the positive input lines that we would directly ground if we used a dual supply. Fig. 7 shows the basic configuration of our circuit, with the 324 pins and power connections drawn in.

Step 4. Choose a frequency range. For CW, let's try 300 to 1200 Hz.

Step 5. Find the center frequency, F_c. This is our first calculation. Let k be the ratio of the highest (F_{hi}) and lowest (F_{lo}) frequencies of our chosen range. Then: $k = F_{hi}/F_{lo} = 1200/300 = 4$.

The square root of k is 2 (and let's call this ks). If we have not made a mistake, then $F_c = F_{hi}/ks = F_{lo}ks = 1200/2 = 300 \times 2 = 600$ Hz.

This is the design center of our filter. Since the filter is tunable, let's next turn to the task of being sure it tunes exactly the range we want it to tune.

Step 6. Choose a dual pot to tune the filter. This is a practical decision. Since we have already said that we would like to keep the feedback resistors well above 10k and hopefully near 100k, a dual 500k pot would be nice. Dual 100k audio-taper pots may be more accessible, so let's see what happens if we use this value.

Step 7. Calculate R_{hi} and R_{lo}. In order to limit the tuning range to specific values (e.g., 300 to 1200 Hz), we will need a fixed resistor and a pot in series to make up each of the frequency-controlling resistors. R_{hi} will be the needed resistance when the frequency is the lowest, and R_{lo} will be the resistance at the highest frequency.

We know something about these values, even though we have not yet selected a capacitor. First, we know that their difference will be 100k, the value of the pot. Hence, $R_{hi} - R_{lo} = 100k$. We also know that $R_{hi} = 4R_{lo}$, since the ratio of low to high frequency is 4:1. (Note: If we hold the capacitance constant, as we will do with a fixed-value unit, the frequency and resistance will vary inversely with each other, i.e., $F_{hi}/F_{lo} = R_{hi}/R_{lo}$.)

Knowing the two relationships between the highest and lowest resistances lets us substitute and solve for R_{lo}. Since the ratio of the resistances is 4:1, then $R_{hi} = 4R_{lo}$. In the difference formula, we now can say that $4R_{lo} - R_{lo} = 100k$, or $3R_{lo} = 100k$. Dividing 100k by 3, we get $R_{lo} = 33.3k$. Since the highest resistance is 100k higher, $R_{hi} = 133.3k$. As a check, we can use the other original formula and let $R_{hi} = 4R_{lo} = 4 \times 33.3k = 133.2k$.

I have carried out the calculation to more precision than we can possibly get with real components to show how good the method is. In fact, since real pots are often shy of 100k by as much as 10%, it is wise to have a pot in hand before working out a design. The decimal places might get long, but rounding to the nearest whole number for resistors and keeping k and ks to no more than two decimal places will give perfectly good design accuracy.

We now know the fixed series resistor for RF1 and RF2 will be 33k, with the 100k pot making up the rest of the resistance. If we discover that our dual pot does not track and can determine by how much it is off, we might make one of the two fixed resistors a 50k trimmer pot. (Adjustment of trimmers in the frequency-determining circuits of a filter is best done with the circuit wired but the op amp out of its socket, using a precise ohmmeter. Accurate adjustment with the circuit in operation requires a scope with frequency-scanning capability. Output-level readings taken on an ac/audio voltmeter can be misleading.)

Step 8. Calculate capacitors CF1 and CF2. At all frequencies, the resistance will equal the capacitive reactance. Hence, the standard formula for calculating capacitance from frequency and reactance becomes $CF1 = CF2 = 1/2piFR_f$. In this case, start with either end of the tuning range. For the example, use 300 Hz, where the resistance is 133k. If your calculator has a 1/X key, you can just multiply all the denominator numbers together and then hit the inverse key. The answer is likely to appear in exponential notation. For example, $C_f = 1/(2 \times 3.14 \times 300 \times 133{,}000) = 3.99 \times 10^{-9}$.

We need to convert this to either microfarads (10^{-6}) or picofarads (10^{-12}) to see what capacitors we should purchase. 3990-pF or .04-uF capacitors will do the job. We can parallel some 5% polystyrene capacitors to hit 4000 pF fairly closely. Given the fact that we can rarely buy the exact value that the formula says we need, we should design the frequency range of the filter with an extra 5% on either end to allow for the slight range shift our approximations will produce.

We can check our work by calculating the two frequency-determining capacitors from the other end of the range. This time, $C_f = 1/(2 \times 3.14 \times 1200 \times 33{,}000) = 4.02 \times 10^{-9}$, or about 4000 pF again. Because we used pi to only two decimal places and dropped the last 300 Ohms off the resistance values, the answers diverge by about 1%, well within the 5% component tolerance. Note that had we used the 500k pot we considered at the beginning of the example, our capacitors would be about one-fifth the present value. Some builders have difficulty obtaining 5% capacitors in the higher values and may want to use the larger pot in order to combine it with capacitors in the 800-pF range.

Step 9. Calculate the resistance at the center frequency, F_c. Since the resistance at center frequency will equal the reactive capacitance, $R_{fc} = 1/2piF_cC_f = 1/(2 \times 3.14 \times 600 \times 4 \times 10^{-9}) = 66{,}348$ Ohms. This is the resistance value of the frequency-determining resistors at the design center of the filter. We will use this figure in a very broad way to determine the remaining resistors in the filter. Most filter-design manuals scale a filter from an initial assumption of equal value resistors throughout as much of the design as possible. On this assumption, R1 through R3 should equal the center-frequency resistance, and R5 should approximate it, if possible. Similar assumptions apply to the other filter designs, with adjustments for values that must differ.

In practice, using components readily accessible to amateurs, the assumption is not very important as long as filter resistor values fall within the range that permits the op amps to perform well. Values from 10k to 100k have been used with no specifically noticeable change of performance. As a rule of thumb, try to let the feedback resistors fall within a 2 to 1 or 3 to 1 ratio of the center-frequency resistance.

Step 10. Determine the feedback and input resistors, R1 through R3. On the basis of the previous calculation and discussion, 68k resistors appear to be the closest value to the calculated center-frequency resistance. In practice, 100k resistors do not change the filter performance. What is important is to use the same value for all three. Since 100k is a nice round value found in most ham junk boxes, let's use it. No-

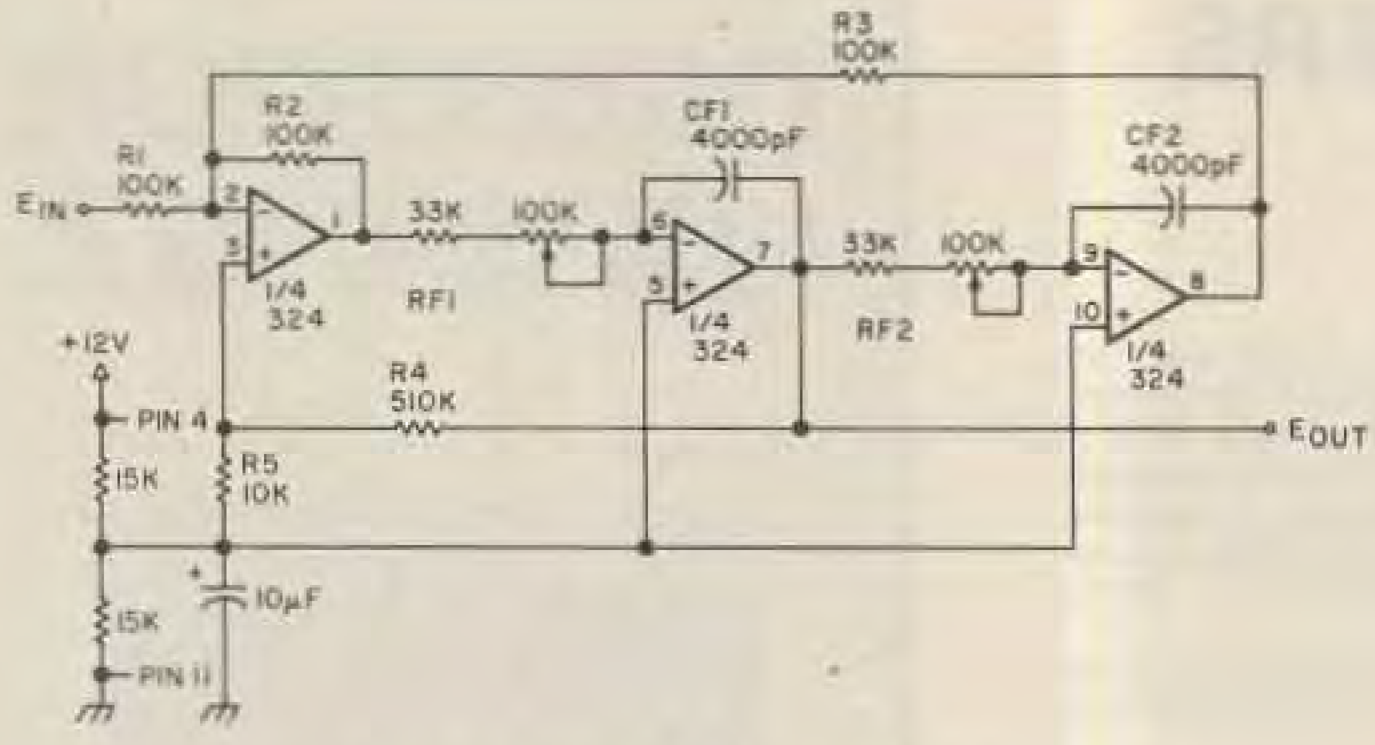

Fig. 8. A fixed-Q −SVF filter.

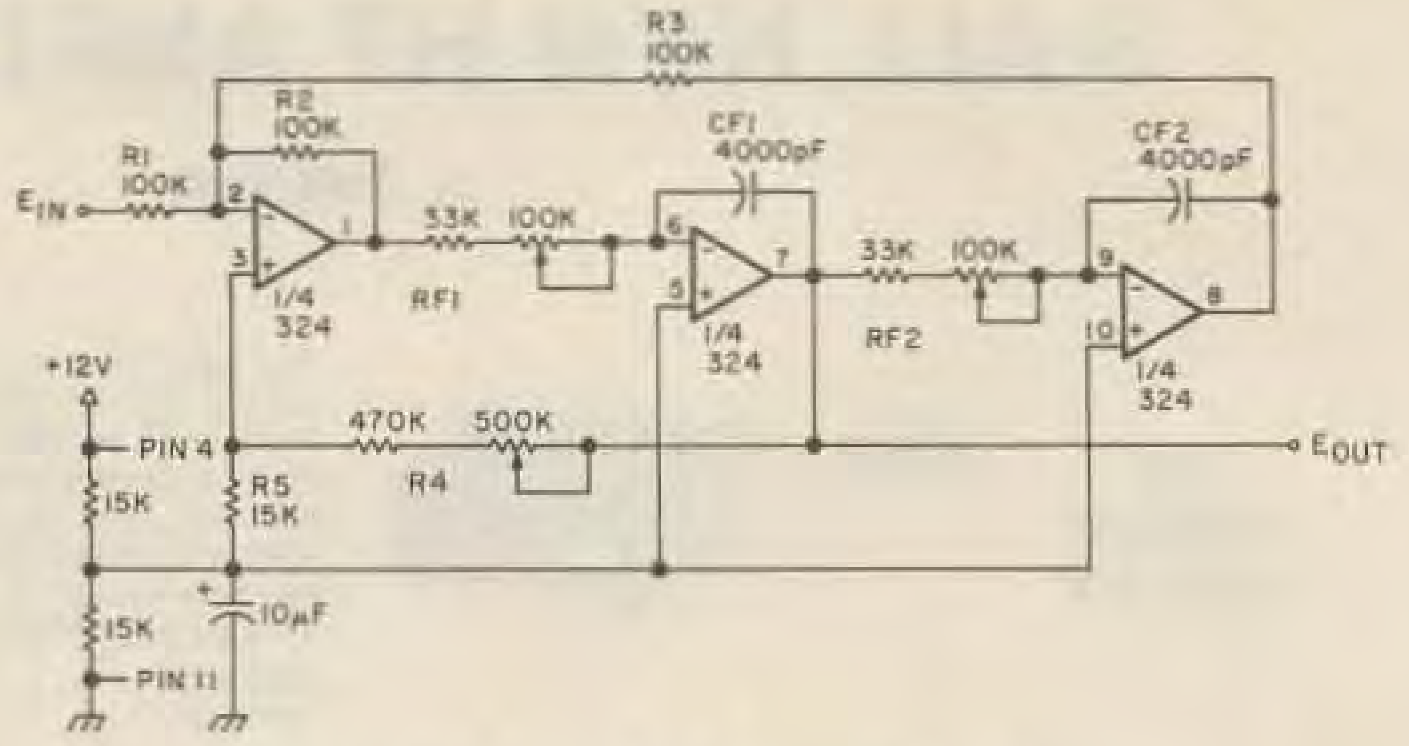

Fig. 9. A variable-Q −SVF filter.

tice that, like many ham building decisions, the grounds for our choice have little relationship to theory. If our value does not work for some reason, we have another value to try.

Step 11. Select a value for Q and choose the Q-determining resistors, R4 and R5. Since both resistors affecting Q and gain are independent of the input resistor, we have more latitude in choosing values than with the other two designs. For CW filters, there is rarely a need for a Q greater than 25, and the range of 10 to 20 will generally produce sufficient selectivity without ringing. For greater selectivity, we should use identical successive filters which will give us a steeper bandwidth curve and greater ultimate rejection on unwanted signals. As a rule of thumb, using 5% and 10% components, I anticipate that the half-voltage (−6 dB) bandwidth will approximate $3F_C/Q$, about 50% wider than theory indicates. For the SVF filters, bandwidth in Hz will vary directly with frequency. Thus, if I choose a 100-Hz bandwidth for the 600-Hz center frequency, it will vary from 50 Hz at the 300-Hz end of the range to 200 Hz at the 1200-Hz upper end of the tuning range. If this bandwidth is acceptable, then $Q = 3F_C/BW_{fc} = (3 \times 600)/100 = 18$. Let's see what happens if we use this figure.

From the formulas governing the −SVF filter, $R4 = R5(3Q-1)$. For our case, $3Q-1 = (3 \times 18) - 1 = 53$, and $R4 = 53R5$. If we let $R4 = 100k$, then $R5 = 5.3$ megohms; use either 4.7-megohm or 5.1-megohm standard resistor values. In fact, we can change the values proportionately by factors of ten without disrupting filter performance. Values of 10k and 510k work well and may be easier to find. A rule of thumb is to let R4 be the highest easy-to-find value that permits R5 (or R1 in the other two designs) to approach its proper theoretic relationship to the other resistors. However, other considerations may enter into the final selection. Fig. 8 shows our completed fixed-Q design.

One major consideration is whether we wish to be able to vary the Q of the filter and thereby to broaden or narrow the bandwidth over some useful range. For example, we might wish to have a Q ranging from 10 to 20 for this design. At $Q = 10$, the resistor ratio $(3Q-1)$ will be 29, and at $Q = 20$, the ratio will be 59. Suppose that we have a 500k pot we wish to use to vary the Q. Since we will not vary the Q to nothing, we will need a series resistor with the pot to make up R4. We know that the value of R4 at $Q = 20$ will be the series resistor R_s + 500k, the highest value of the pot. At $Q = 10$, R4 will be just R_s, the value of the fixed series resistor. At the higher Q, $R5 = (R_s + 500{,}000)/59$, while at the lower Q, $R5 = R_s/29$. We can solve for the series resistor by letting $R_s/29 = (R_s + 500{,}000)/59$. Cross multiplying, we get $30R_s = 29 \times 500{,}000$, or $R_s = 1{,}450{,}000/30 = 483{,}333$ Ohms. This is the series resistor to go with the 500k pot for R4. $R5 = R_s/29 = 483{,}333/29 = 16{,}667$ Ohms. (As a check, $R5 = (483{,}333 + 500{,}000)/59 = 16{,}667$.) We can choose a 15k or 18k resistor for R5 and a 470k or 510k resistor for R_s, respectively. Exactness will not matter too much here since we will tune the control for best reception rather than for some specific value of Q. Fig. 9 displays our completed variable-Q design.

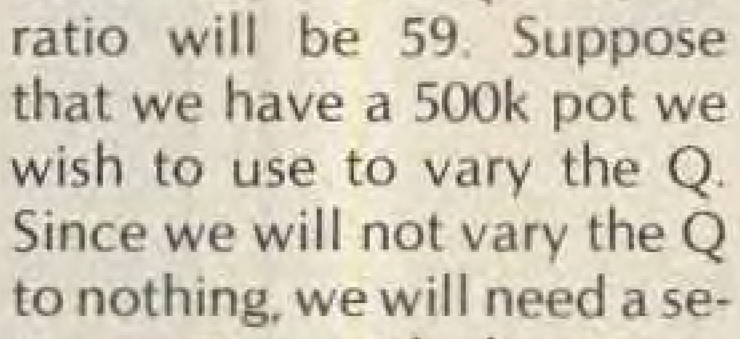

Fig. 10. Filter input- and output-level controls.

Step 12. Consider the gain. This step does not require special calculations, but it does bring the matter of gain to your attention. For the −SVF design, gain will equal Q. If you design a fixed-Q filter, you can accommodate the filter gain with preceding and succeeding level controls, as shown in Fig. 10. Set the input-level control so that the strongest signal will not drive the filter to clipping. A scope will show this as a sharply flattened sine wave. Since the voltage gain will be considerable, the filter may drive the succeeding stage too hard, causing distortion in the amplifier feeding the speaker or phones. We can kill the unwanted voltage with another trimmer set to hold the amplifier relatively distortionless at full volume.

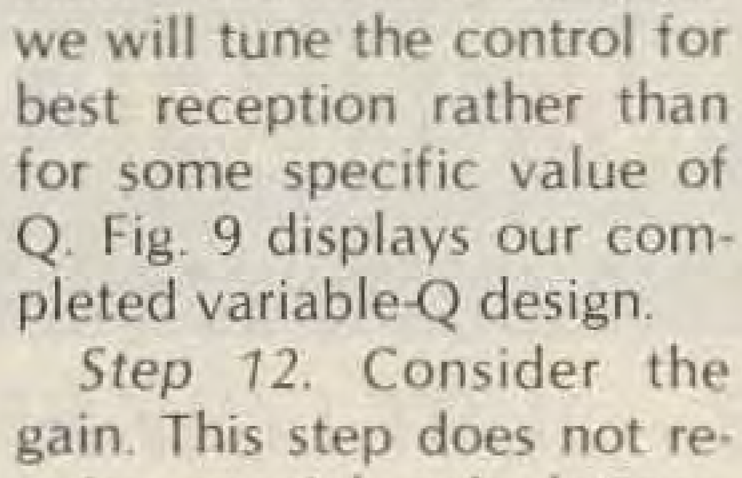

Fig. 11. A simple post-filter amplifier for speaker or phones.

If the filter has a variable-Q control, then its gain will also vary. To avoid the need for constant volume-control adjustments, the compression amplifier shown in Fig. 4 should follow the filter and precede the output amplifier. With the values shown for the compression circuit, a normal CW signal will leave the speaker quiet between dots and dashes. The circuit needs no input-setting pot, and the output-level control serves the same function as the filter-output control in Fig. 10.

These 12 steps complete the design phase of the work. The next step is to breadboard a model, verify its operation, and finally construct a permanent version complete with case and power source. Robbing power from the receiver and installing the filter in either the receiver cabinet (especially if inserted between the detector and audio stages) or the speaker cabinet (along with an audio amplifier such as the LM386 circuit shown in Fig. 11) is one popular way to handle final construction. However, to avoid cabinet and circuit modifications, you may wish to make the filter a self-contained unit.

Additional Procedures— +SVF and B-Q Designs

The first eight steps of the procedures just outlined are identical for all three filter designs. Nothing changes until Step 10, selection of the input resistor, where we have only a minor modification for the +SVF filter. R1 should be half the value of either R2 or R3 if we wish to have the relationship of Q and gain follow the formulas given with Fig. 5. Other ratios are possible, although the input resistor should not be greater than the feedback resistors. The gain will change but remain constant across the tuning range.

Let's look more closely at the final steps of the procedure, customizing them for each particular design. First, the bi-quad filter:

Step 11: B-Q. Select a value for Q, and choose the Q-determining resistors. In the B-Q design, the input resistor, R1, interacts with R4 to determine Q and gain. Having selected an input resistor, R4 = QR1. Selecting Q follows the same guidelines given for the −SVF design, with the proviso that Q will vary across the tuning range, since bandwidth in Hertz is constant. Using our −6-dB (half-voltage point) rule of thumb, we can design with the formula $Q = 3F_c/BW$, where BW is the desired bandwidth in Hertz. If we wish about 100 Hz, then Q = (3 × 600)/100 = 18. RV = 18R1 = 18 × 100k = 1.8 megohms, a usable value. However, with very little change in performance, we can reduce both R1 and R4 as long as we keep them in the proper ratio. Fig. 12 shows the full results of our design work.

We can vary the Q and consequently the bandpass of B-Q filters. We need only make R4 variable. Suppose we wish to vary the Q between about 10 and 20. If R1 is 100k, then R4 needs to be 1 megohm for a Q of 10 and 2 megohms for a Q of 20. We can use a one-meg fixed resistor in series with a one-meg pot for R4, and the problem is solved. Fig. 13 shows the changes necessary for variable Q.

Step 12: B-Q. Consider the gain. Variable Q plus the natural gain variability of the B-Q filter makes a compression amplifier almost mandatory. However, the 100-to-1 compression capability of the audio elixir circuit will more than cover the situation. The natural gain variability of a fixed B-Q filter with the 300-to-1200-Hz tuning range is about 4 to 1, while Q variability expands the total range to 40 to 1, well within the amplifier's capabilities and with room to spare for audio signal strength variations.

The B-Q filter has one special property not shared by either SVF design. You can tune the B-Q using only RF1, leaving RF2 fixed for F_c. The variable resistor, however, will change frequency only with the square root of the resistance change, meaning that the pot will have to have a much wider range to cover the chosen frequency range. Since the frequency limits in the example are $2F_c$ and $F_c/2$, the resistance range must be $R_{fc}/4$ and $4R_{fc}$. In this design, $R_{fc} = 66{,}348$ Ohms. The lowest resistance (for the highest frequency) will be 66,348/4 = 16,587, while the highest resistance (for the lowest frequency) will be 66,348 × 4 = 265,392. The difference is 248,805. A 250k pot in series with a 15k fixed resistor will form a satisfactory RF1. An audio taper or reverse log pot is mandatory in this application, since even with a log pot the frequency will compress at one end of the scale.

In this example, we were fortunate to wind up with a required value close to an existing potentiometer value. For designing a single pot B-Q filter from scratch, we can begin at Step 6, choosing a pot to tune the filter. Let's select a 500k pot and see what happens.

Step 7: B-Q, single pot. Calculate R_{hi} and R_{lo}. Since frequency will vary as the square root of resistance changes, the total resistance change will be k^2, where k is the frequency ratio. Since k = 4 (1200/300 Hz), $k^2 = 16$. $R_{hi} = 16R_{lo}$. We also know that $R_{hi} = R_{lo} + 500k$. Now we can solve for R_{lo}: $16R_{lo} = R_{lo} + 500{,}000$, or $R_{lo} = 500{,}000/15 = 33{,}333$ Ohms. This is the value of the fixed-series resistor. $R_{hi} = 33{,}333 + 500{,}000 = 533{,}333$ Ohms. As a check, 533,333/16 = 33,333 Ohms.

The resistance at center frequency (and fixed frequency-determining resistor RF2) will be $R_{hi}/4 = 4R_{lo} = 533{,}333/4 = 33{,}333 \times 4 = 133{,}333$ Ohms. We can use 100k and 33k resistors in series or use the nearest standard value.

Step 8: B-Q, single pot. Calculate capacitors CF1 and CF2. This calculation uses the same procedure as in the −SVF filter. Since resistance and capacitive reactance are the same at the center frequency (and we must use F_c for this calculation), $C_f = 1/2piF_cRF2 = 1/(2 \times 3.14 \times 600 \times 133{,}333) = 1.99 \times 10^{-9}$. This is about 2000 pF, an obtainable value in polystyrene capacitors.

Determine the remaining values for the filter in the ordinary way. 100k feedback and input resistors appear to be in order, since they vary only a little from the value of RF2. Considerations of Q and gain will be identical to those for the dual-pot biquad design. Fig. 14 shows our new filter.

The SVF filters always require dual pots. Therefore, the only difference between the +SVF filter and the −SVF design concerns Q and gain.

Step 11: +SVF. Select a value for Q, and choose the Q-determining resistors. Q selection for the +SVF is

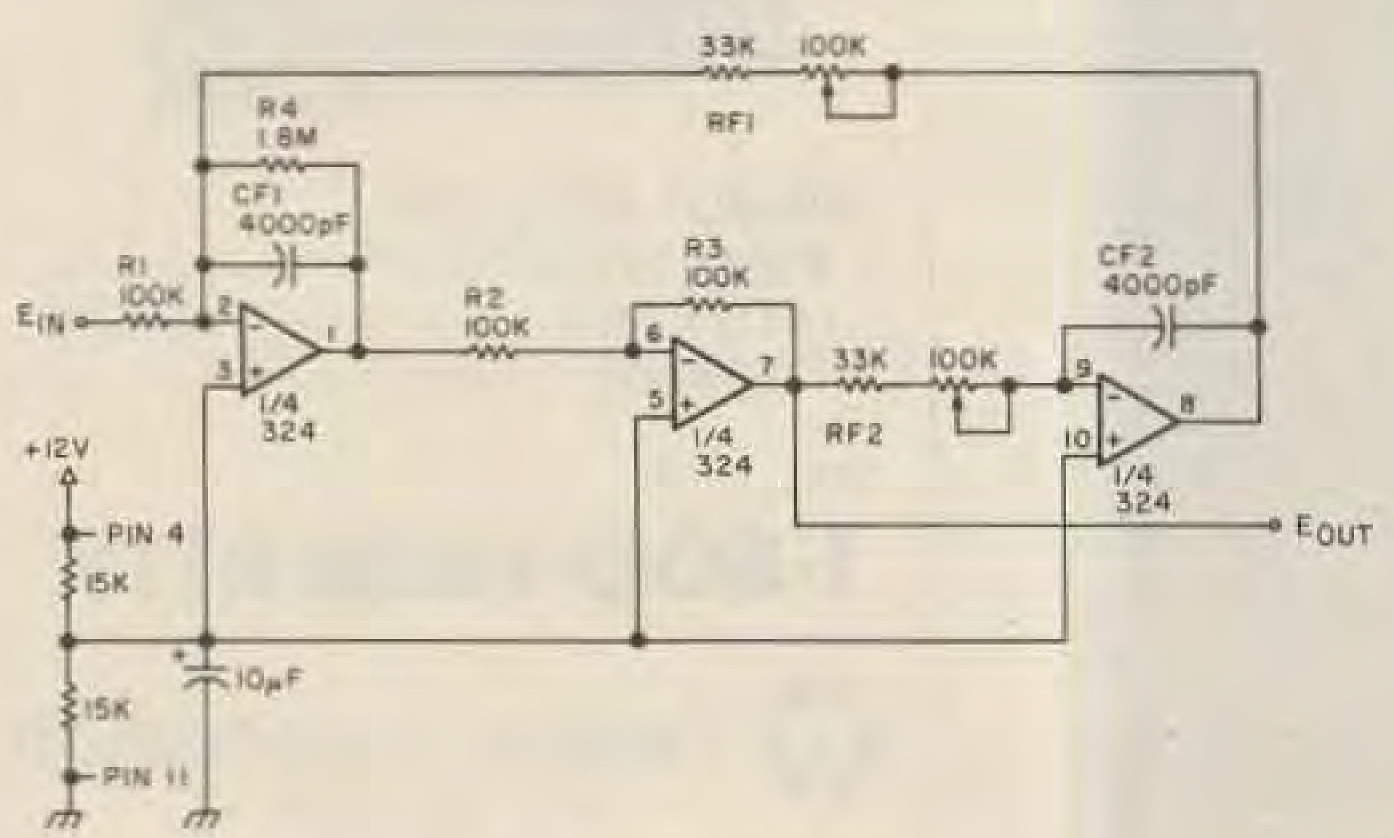

Fig. 12. A fixed-Q B-Q filter.

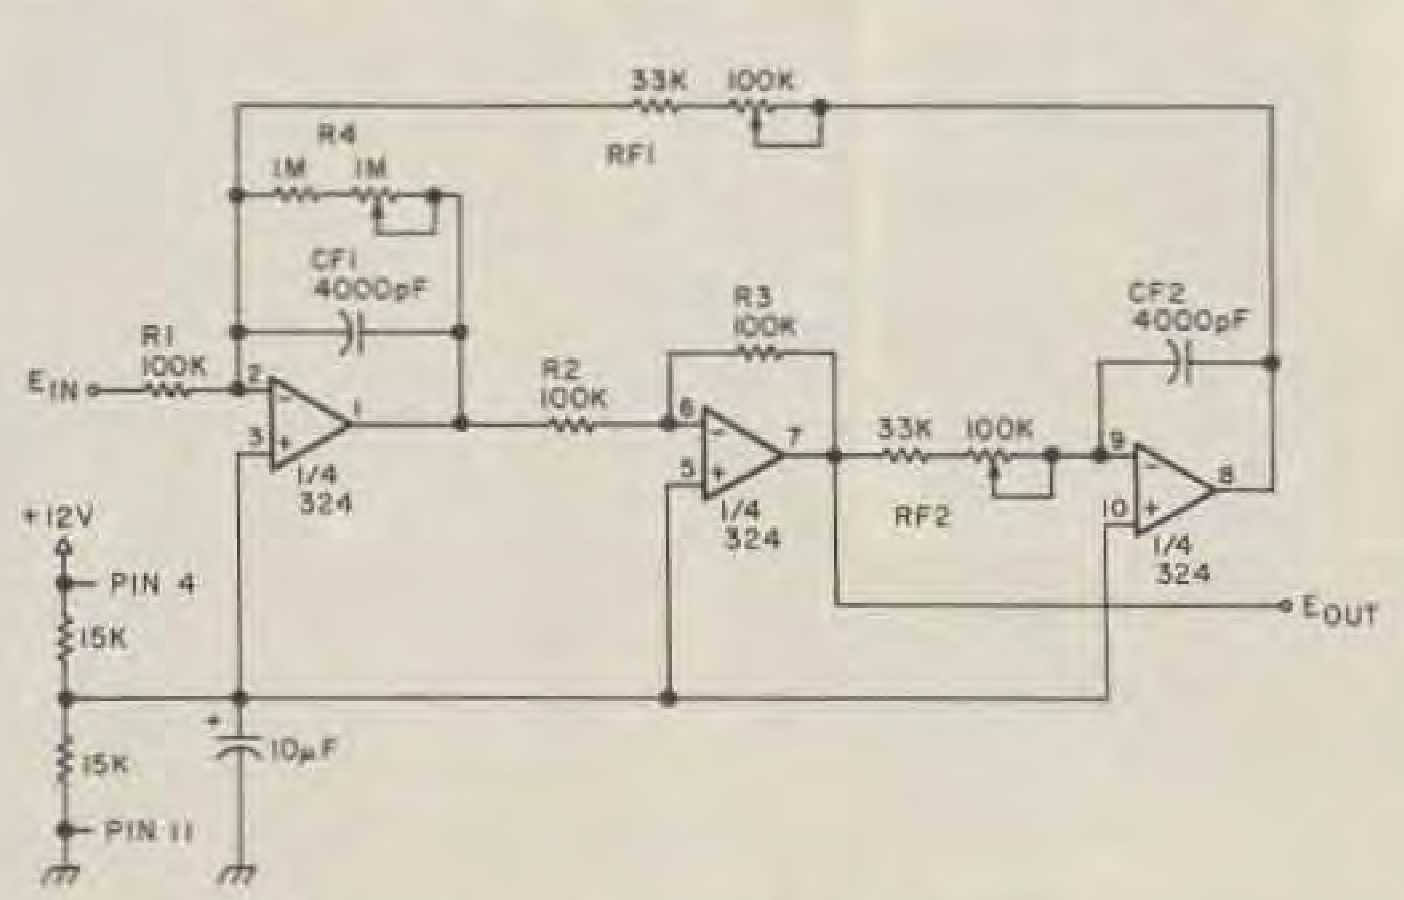

Fig. 13. A variable-Q B-Q filter.

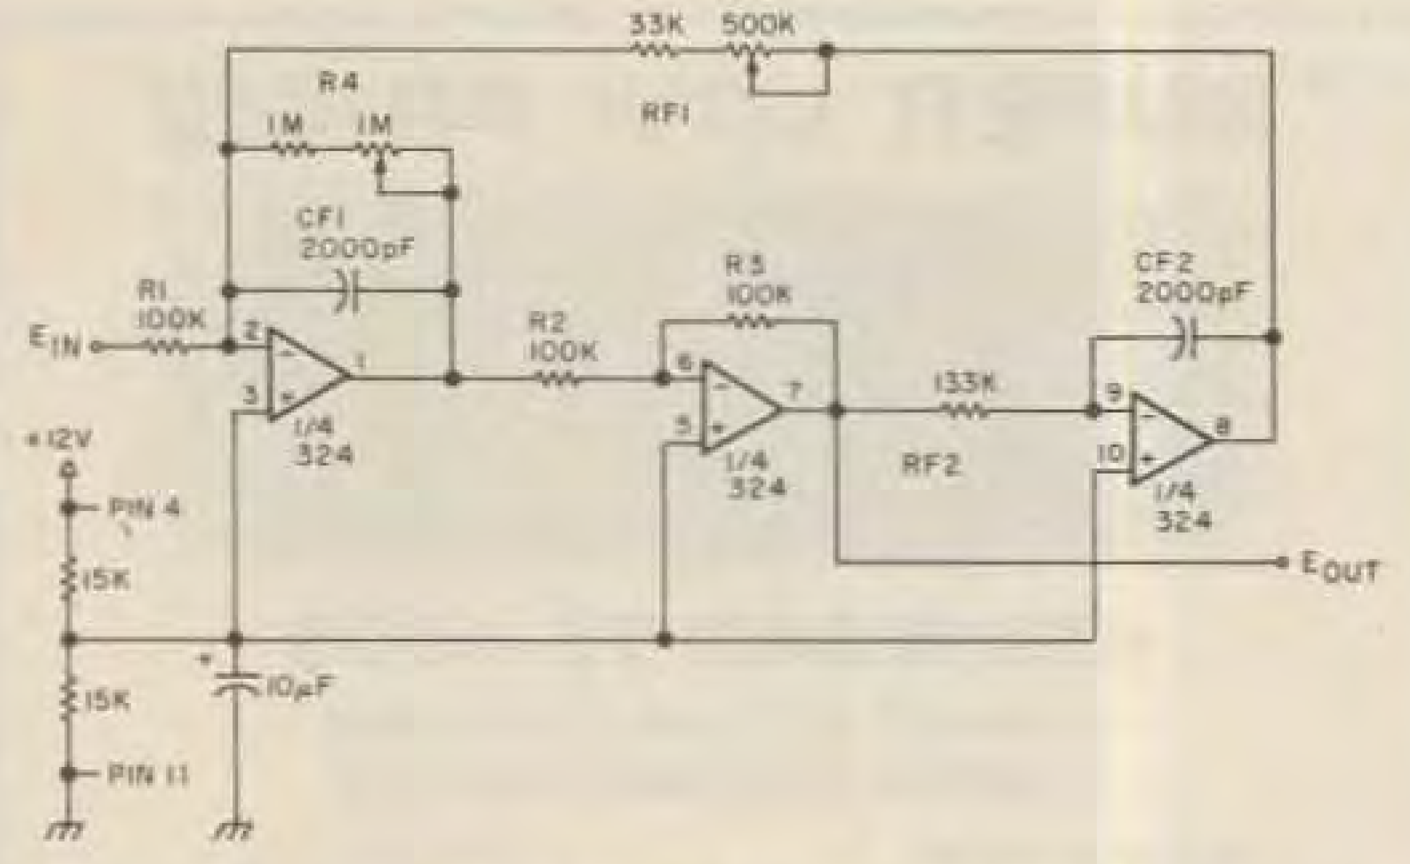

Fig. 14. A single-pot tunable variable-Q B-Q filter.

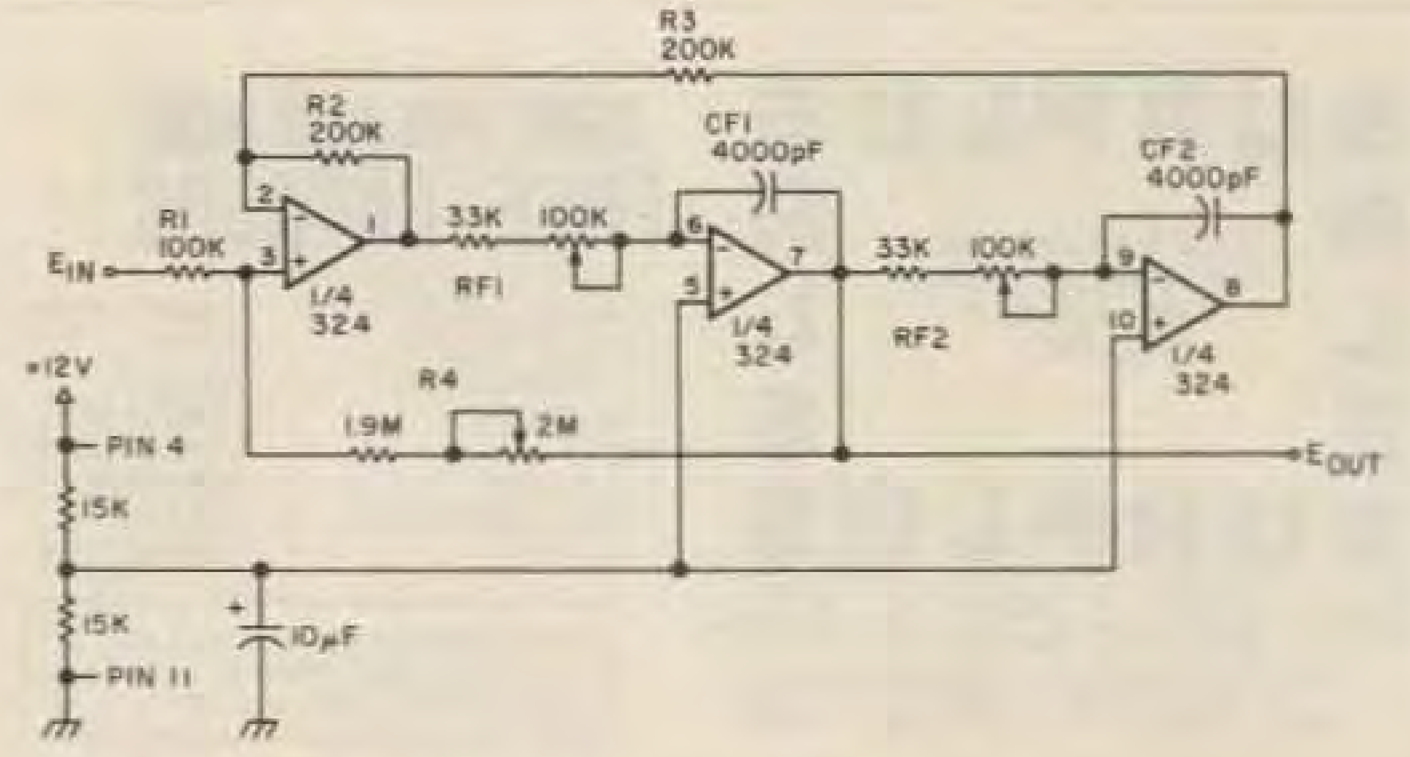

Fig. 15. A variable-Q +SVF filter.

identical to that for the −SVF design. We must make mental note that gain will double Q if we follow recommended resistor relationships. Let Q = 18. R4 = R1 (2Q − 1). If we use 100k resistors for feedback, the R1 is 50k. Many designs use 200k values for R2 and R3, in which case, R1 = 100k. Let's use this latter value for our design. For a Q of 18, 2Q − 1 = 35, and therefore R4 = 100k × 35 = 3.5 megohms. 3.3 megohms would work well. For a variable Q of, say, 10 to 20, the maximum resistance value of R4 would be 39R1 and the minimum value would be 19R1. R4 will range from a series resistor value of R_s to R_s + pot, where pot is the potentiometer value we select. Let's try a 2-megohm pot. Then $R1 = R_s/19$ at low Q and $(R_s + 2{,}000{,}000)/39$ at high Q. Solving for R_s, we get $R_s = 38{,}000{,}000/20 = 1.9$ megohms. R4 thus becomes a 1.9-megohm fixed resistor in series with a 2-megohm pot. $R1 = R4/(2Q - 1) = 3.9$ megohms/39 = 1.9 megohms/19 = 100k, a desirable value.

Step 12: +SVF. Consider the gain. The gain of this +SVF filter, shown in Fig. 15, ranges from 20 to 40, depending upon the variable Q. Again, following this design with a compression amplifier is a must for easy use.

Construction and Results

All of the designs shown in the examples have been breadboarded to confirm that they will work. In fact, they all work even when some non-frequency-determining components vary by 20% from the design values. Fig. 16 charts the test results. (Always test a design on a breadboard before wiring a final version. If nothing else, the breadboard test will turn up bad components. More important, adjusting the design to more precisely meet your needs is much simpler on a breadboard.)

Construction of the final model can take any form. Perfboard and printed circuit board peform equally well. Layout is not critical with the LM324. The TL084 requires some care to prevent inadvertent coupling, a more serious concern with the very high impedance inputs to each section. One easy way to overcome the problem is to avoid compressing the components into too small a space. Spreading the fixed components at the IC corners in a radial pattern tends to prevent unwanted coupling and makes component replacement simpler. Beyond this, construction is left to individual ingenuity.

Part of the construction ease stems from the low Q of these filters. Most practical filter articles still manage to repeat the virtually useless fact that these designs are good to a Q of 500. At normal CW audio, the bandwidth would be just over 1 Hz, and the filter would ring for a week with just one receiver electron pop, if it was not already oscillating. With normal components, practical Qs of 5 to 25 ensure good stability and

Figure	Filter	Tuning Range	Bandwidth	Q	Output Voltage Ratio	Notes
8	−SVF, fixed Q	330–1250 Hz	25–100 Hz	26	1.08:1 F_{hi}:F_{lo}	3900 pF, 5% C*; reduce R4 from 510k to 330k**
9	−SVF, variable Q	330–1250 Hz	40–75 Hz at (at F_c)	16 to 30	2.5:1 Q_{hi}:Q_{lo}	Reduce fixed-series Q resistor from 470k to 330k
12	B-Q, fixed Q	330–1250 Hz	30 Hz	35 at F_C	1.4:1 F_{hi}:F_{lo}	Reduce R4 from 1.8 megohms to <820k
13	B-Q, variable Q	330–1250 Hz	30 Hz at Q_{hi} 50 Hz at Q_{lo} (at F_c)	35 21	2:1 Q_{hi}:Q_{lo} 2.8:1 total change due to both F and Q	Reduce fixed series Q resistor from 1 megohm to 470k
14	B-Q, single pot, variable Q	350–1330 Hz	55 Hz at Q_{hi} 140 Hz at Q_{lo} (at F_c)	18 7.3	2.6:1 total change due to both F and Q	2000 pF, 10% C_f; increase fixed tuning resistor from 33k to >47k to adjust tuning range
15	+SVF, variable Q	340–1260 Hz	40 Hz at Q_{hi} 75 Hz at Q_{lo} (in passband)	32 17	1.6:1 total change due to both F and Q at R4	Reduce fixed series Q resistor from 1.9 megohms to <1 megohm

Notes: *All filters except the single-pot B-Q used 3900-pF, 5% polystyrene capacitors. The single-pot B-Q model used 2000-pF, 10% polystyrene capacitors. **In any of the filters, raise or lower Q by raising or lowering R4, the Q-determining resistor.

Fig. 16. Test results and comments on the six sample filters.

easy adjustment for ham audio filters. Using 5% and 10% tolerance components (or bridge-matched 20% components), Q will be slightly less than theory predicts but more than adequate. Moreover, a close examination of the peak of the response curve will reveal irregularities compared to the theoretic ideal, but these will always be too small to make a difference to CW or to phone reception. In short, for all practical purposes, home brew can be as good as commercial when it comes to simple audio filters to enhance reception. To the degree that we can customize the design to our specific needs, they might even be better than commercial for some hams.

The purpose of outlining these procedures is to reduce the design of custom ham CW bandpass filters to a series of steps that ensures not only a filter that works, but, as well, a filter that tunes frequency and Q over just the operator's desired range. The procedures are applicable to today's run of multiple ICs such as the LM324, the TL084, and the 3900 Norton. As new generations of op amps emerge, with different biasing, input, and current requirements, the rules of thumb will likely change. However, the basic principles of determining RC tuning ranges will not. Only our selection of pots and fixed resistors will vary for new impedance-matching conditions. Hence, with adaptation for new devices, CW operators and other hams who need bandpass filters using the usual lot of reasonably priced 5% and 10% components should be able to satisfy their needs on their own work benches. A hand calculator, a sharp pencil, and a breadboard are the basic tools for good filtering. Triple op-amp designs just make the task a bit more challenging. ■