THE ART OF ELECTRONICS

Second Edition

Paul Horowitz HARVARD UNIVERSITY

Winfield Hill ROWLAND INSTITUTE FOR SCIENCE, CAMBRIDGE, MASSACHUSETTS

CAMBRIDGE
UNIVERSITY PRESS

ACTIVE FILTERS AND OSCILLATORS

CHAPTER 5

With only the techniques of transistors and op-amps it is possible to delve into a number of interesting areas of linear (as contrasted with digital) circuitry. We believe that it is important to spend some time doing this now, in order to strengthen your understanding of some of these difficult concepts (transistor behavior, feedback, op-amp limitations, etc.) before introducing more new devices and techniques and getting into the large area of digital electronics. In this chapter, therefore, we will treat briefly the areas of active filters and oscillators. Additional analog techniques are treated in Chapter 6 (voltage regulators and high-current design), Chapter 7 (precision circuits and low noise), Chapter 13 (radiofrequency techniques), Chapter 14 (low-power design), and Chapter 15 (measurements and signal processing). The first part of this chapter (active filters, Sections 5.01–5.11) describes techniques of a somewhat specialized nature, and it can be passed over in a first reading. However, the latter part of this chapter (oscillators, Sections 5.12–5.19) describes techniques of broad utility and should not be omitted.

ACTIVE FILTERS

In Chapter 1 we began a discussion of filters made from resistors and capacitors. Those simple *RC* filters produced gentle high-pass or low-pass gain characteristics, with a 6dB/octave falloff well beyond the −3dB point. By cascading high-pass and low-pass filters, we showed how to obtain bandpass filters, again with gentle 6dB/octave "skirts." Such filters are sufficient for many purposes, especially if the signal being rejected by the filter is far removed in frequency from the desired signal passband. Some examples are bypassing of radiofrequency signals in audio circuits, "blocking" capacitors for elimination of dc levels, and separation of modulation from a communications "carrier" (see Chapter 13).

5.01 Frequency response with *RC* filters

Often, however, filters with flatter passbands and steeper skirts are needed. This happens whenever signals must be filtered from other interfering signals nearby in frequency. The obvious next question is whether or not (by cascading a number of identical low-pass filters, say) we can generate an approximation to the ideal "brick-wall" low-pass frequency response, as in Figure 5.1.

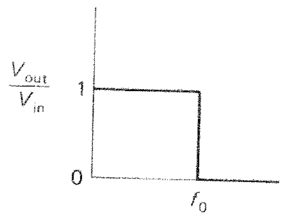

Figure 5.1

We know already that simple cascading won't work, since each section's input impedance will load the previous section seriously, degrading the response. But with buffers between each section (or by arranging to have each section of much higher impedance than the one preceding it), it would seem possible. Nonetheless, the answer is no. Cascaded *RC* filters do produce a steep *ultimate* falloff, but the "knee" of the curve of response versus frequency is not sharpened. We might restate this as "many soft knees do not a hard knee make." To make the point graphically, we have plotted some graphs of gain response (i.e., V_{out}/V_{in}) versus frequency for low-pass filters constructed from 1, 2, 4, 8, 16, and 32 identical *RC* sections, perfectly buffered (Fig. 5.2).

The first graph shows the effect of cascading several *RC* sections, each with its 3dB point at unit frequency. As more sections are added, the overall 3dB point is pushed downward in frequency, as you could easily have predicted. To compare filter characteristics fairly, the rolloff frequencies of the individual sections should be adjusted so that the overall 3dB point is always at the same frequency. The other graphs in Figure 5.2, as well as the next few graphs in this chapter, are all "normalized" in frequency, meaning that the −3dB point

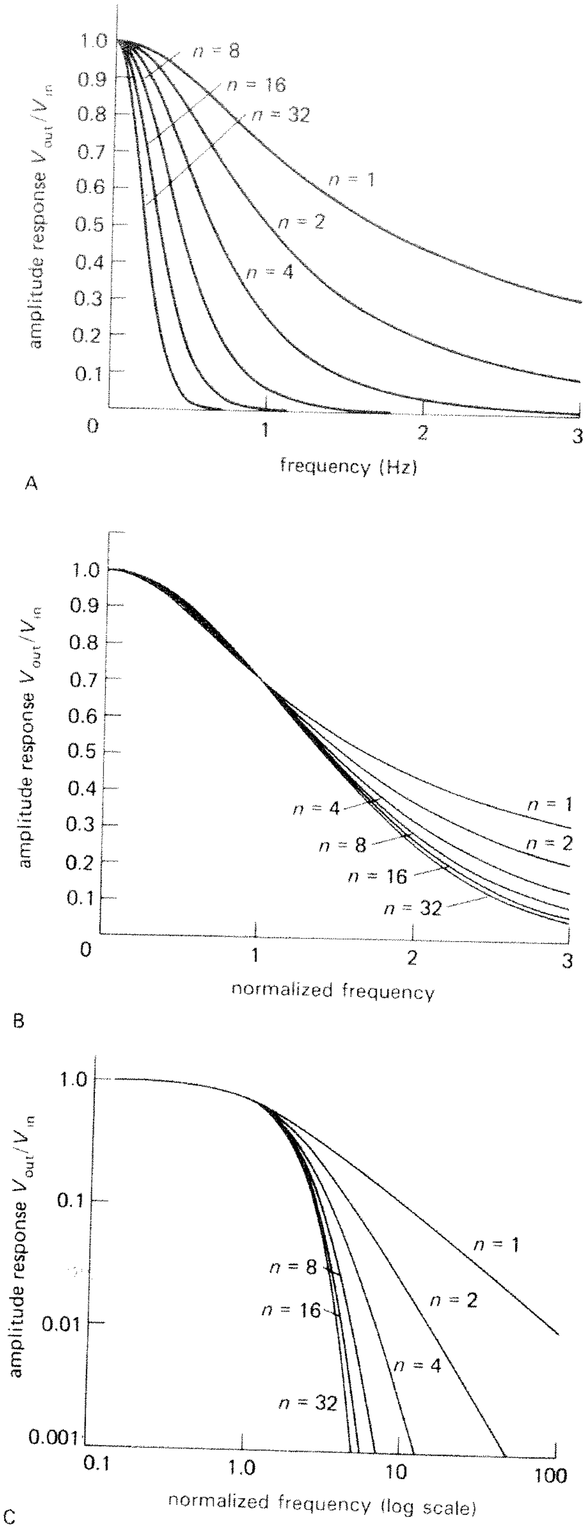

Figure 5.2. Frequency responses of multisection *RC* filters. Graphs A and B are linear plots, whereas C is logarithmic. The filter responses in B and C have been normalized (or scaled) for 3dB attenuation at unit frequency.

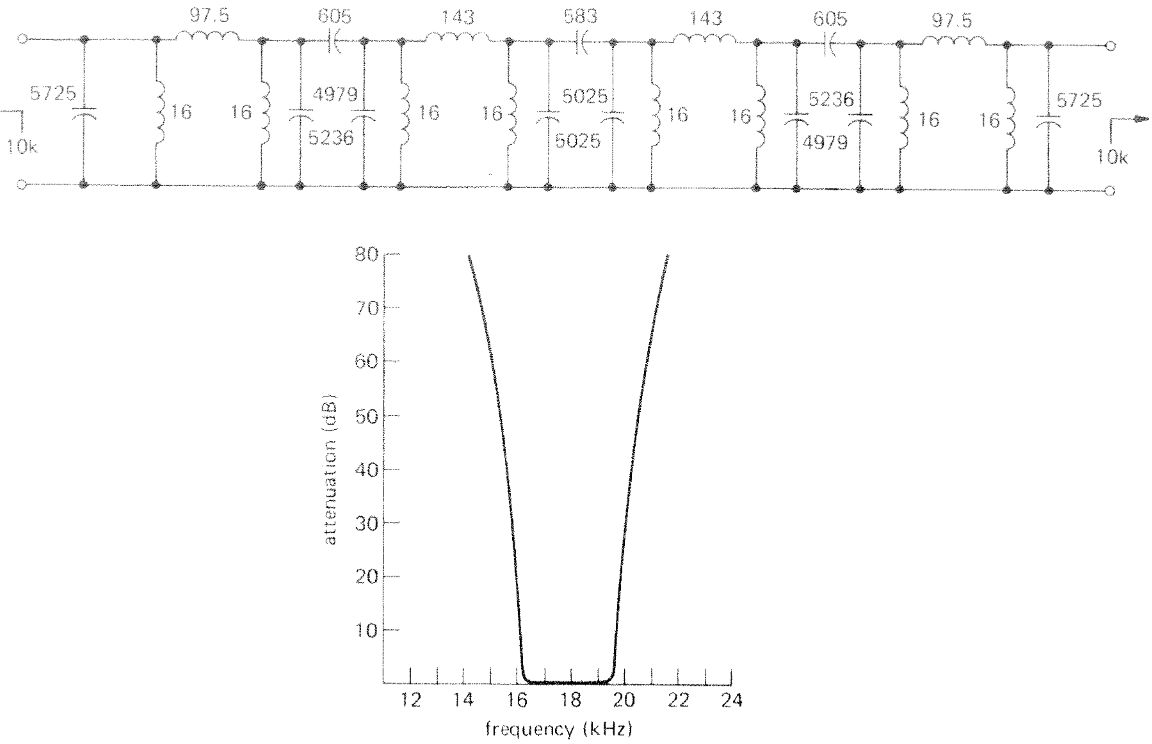

Figure 5.3. An unusually good passive bandpass filter implemented from inductors and capacitors (inductances in mH, capacitances in pF). Bottom: Measured response of the filter circuit. [Based on Figs. 11 and 12 from Orchard, H. J., and Sheahan, D. F., *IEEE Journal of Solid-State Circuits*, Vol. SC-5, No. 3 (1970).]

(or breakpoint, however defined) is at a frequency of 1 radian per second (or at 1Hz). To determine the response of a filter whose breakpoint is set at some other frequency, simply multiply the values on the frequency axis by the actual breakpoint frequency f_c. In general, we will also stick to the log-log graph of frequency response when talking about filters, because it tells the most about the frequency response. It lets you see the approach to the ultimate rolloff slope, and it permits you to read off accurate values of attenuation. In this case (cascaded *RC* sections) the normalized graphs in Figures 5.2B and 5.2C demonstrate the soft knee characteristic of passive *RC* filters.

5.02 Ideal performance with *LC* filters

As we pointed out in Chapter 1, filters made with inductors and capacitors can have very sharp responses. The parallel *LC* resonant circuit is an example. By including inductors in the design, it is possible to create filters with any desired flatness of passband combined with sharpness of transition and steepness of falloff outside the band. Figure 5.3 shows an example of a telephone filter and its characteristics.

Obviously the inclusion of inductors into the design brings about some magic that cannot be performed without them. In the terminology of network analysis, that magic consists in the use of "off-axis poles." Even so, the complexity of the filter increases according to the required flatness of passband and steepness of falloff outside the band, accounting for the large number of components used in the preceding filter. The transient response and phase-shift characteristics are also generally degraded as the amplitude response is improved to approach the ideal brick-wall characteristic.

The synthesis of filters from passive components (*R, L, C*) is a highly developed subject, as typified by the authoritative handbook by Zverev (see chapter references at end of book). The only problem is that inductors as circuit elements frequently leave much to be desired. They are often bulky and expensive, and they depart from the ideal by being "lossy," i.e., by having significant series resistance, as well as other "pathologies" such as nonlinearity, distributed winding capacitance, and susceptibility to magnetic pickup of interference.

What is needed is a way to make inductorless filters with the characteristics of ideal *RLC* filters.

5.03 Enter active filters: an overview

By using op-amps as part of the filter design, it is possible to synthesize any *RLC* filter characteristic without using inductors. Such inductorless filters are known as active filters because of the inclusion of an active element (the amplifier).

Active filters can be used to make low-pass, high-pass, bandpass, and band-reject filters, with a choice of filter types according to the important features of the response, e.g., maximal flatness of passband, steepness of skirts, or uniformity of time delay versus frequency (more on this shortly). In addition, "all-pass filters" with flat amplitude response but tailored phase versus frequency can be made (they're also known as "delay equalizers"), as well as the opposite – a filter with constant phase shift but tailored amplitude response.

☐ *Negative-impedance converters and gyrators*

Two interesting circuit elements that should be mentioned in any overview are the negative-impedance converter (NIC) and the gyrator. These devices can mimic the properties of inductors, while using only resistors and capacitors in addition to op-amps.

Once you can do that, you can build inductorless filters with the ideal properties of any *RLC* filter, thus providing at least one way to make active filters.

The NIC converts an impedance to its *negative*, whereas the gyrator converts an impedance to its *inverse*. The following exercises will help you discover for yourself how that works out.

EXERCISE 5.1

Show that the circuit in Figure 5.4 is a negative-impedance converter, in particular that $Z_{in} = -Z$. Hint: Apply some input voltage V, and compute the input current I. Then take the ratio to find $Z_{in} = V/I$.

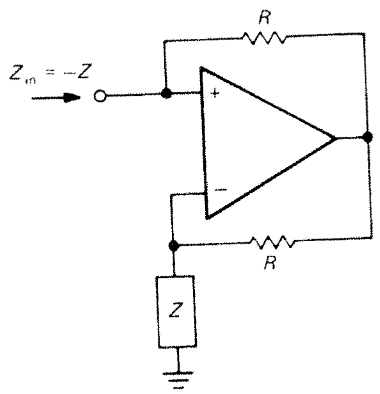

Figure 5.4. Negative-impedance converter.

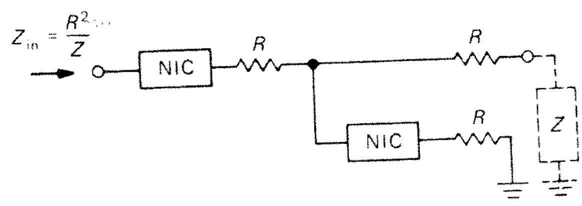

Figure 5.5

EXERCISE 5.2

Show that the circuit in Figure 5.5 is a gyrator, in particular that $Z_{in} = R^2/Z$. Hint: You can analyze it as a set of voltage dividers, beginning at the right.

The NIC therefore converts a capacitor to a "backward" inductor:

$$Z_C = 1/j\omega C \rightarrow Z_{\text{in}} = j/\omega C$$

i.e., it is inductive in the sense of generating a current that lags the applied voltage, but its impedance has the wrong frequency dependence (it goes down, instead of up, with increasing frequency). The gyrator, on the other hand, converts a capacitor to a true inductor:

$$Z_C = 1/j\omega C \rightarrow Z_{\text{in}} = j\omega CR^2$$

i.e., an inductor with inductance $L = CR^2$.

The existence of the gyrator makes it intuitively reasonable that inductorless filters can be built to mimic any filter using inductors: Simply replace each inductor by a gyrated capacitor. The use of gyrators in just that manner is perfectly OK, and in fact the telephone filter illustrated previously was built that way. In addition to simple gyrator substitution into preexisting RLC designs, it is possible to synthesize many other filter configurations. The field of inductorless filter design is extremely active, with new designs appearing in the journals every month.

Sallen-and-Key filter

Figure 5.6 shows an example of a simple and even partly intuitive filter. It is known as a Sallen-and-Key filter, after its inventors. The unity-gain amplifier can be an op-amp connected as a follower, or just an emitter follower. This particular filter is a 2-pole high-pass filter. Note that it would be simply two cascaded RC high-pass filters except for the fact that the bottom of the first resistor is bootstrapped by the output. It is easy to see that at very low frequencies it falls off just like a cascaded RC, since the output is essentially zero. As the output rises at increasing frequency, however, the bootstrap action tends to reduce the attenuation, giving a sharper knee. Of course, such hand-waving cannot substitute for honest analysis, which luckily has already been done for a prodigious variety of nice filters. We will come back to active filter circuits in Section 5.06.

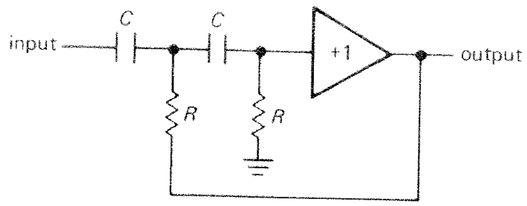

Figure 5.6

5.04 Key filter performance criteria

There are some standard terms that keep appearing when we talk about filters and try to specify their performance. It is worth getting it all straight at the beginning.

Frequency domain

The most obvious characteristic of a filter is its gain versus frequency, typified by the sort of low-pass characteristic shown in Figure 5.7.

The *passband* is the region of frequencies that are relatively unattenuated by the filter. Most often the passband is considered to extend to the −3dB point, but with certain filters (most notably the "equiripple" types) the end of the passband may be defined somewhat differently. Within the passband the response may show variations or *ripples*, defining a *ripple band*, as shown. The *cutoff frequency*, f_c, is the end of the passband. The response of the filter then drops off through a *transition region* (also colorfully known as the *skirt* of the filter's response) to a *stopband*, the region of significant attenuation. The stopband may be defined by some minimum attenuation, e.g., 40dB.

Along with the gain response, the other parameter of importance in the frequency

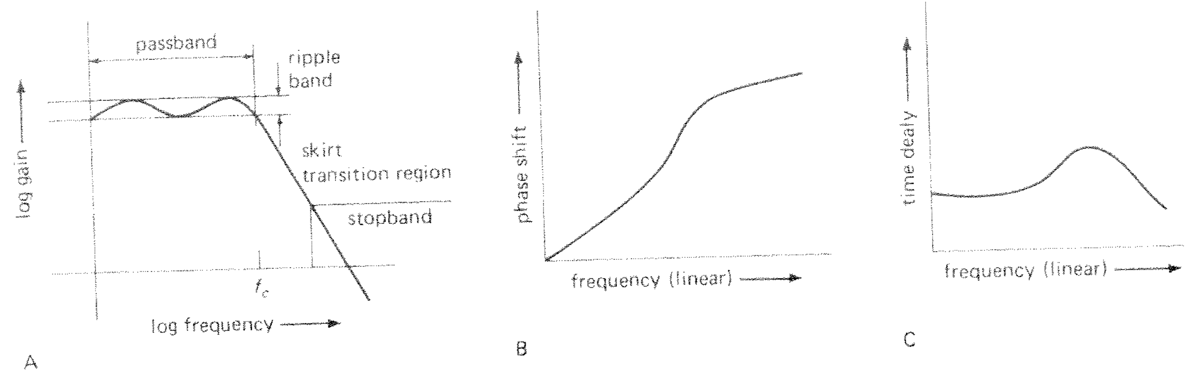

Figure 5.7. Filter characteristics versus frequency.

domain is the *phase shift* of the output signal relative to the input signal. In other words, we are interested in the *complex* response of the filter, which usually goes by the name of **H**(s), where s = $j\omega$, where **H**, s, and ω all are complex. Phase is important because a signal entirely within the passband of a filter will emerge with its waveform distorted if the time delay of different frequencies in going through the filter is not constant. Constant time delay corresponds to a phase shift increasing linearly with frequency; hence the term *linear-phase filter* applied to a filter ideal in this respect. Figure 5.8 shows a typical graph of phase shift and amplitude for a low-pass filter that is definitely not a linear-phase filter. Graphs of phase shift versus frequency are best plotted on a linear-frequency axis.

Time domain

As with any ac circuit, filters can be described in terms of their *time-domain* properties: rise time, overshoot, ringing, and settling time. This is of particular importance where steps or pulses may be used. Figure 5.9 shows a typical low-pass-filter step response. Here, *rise time* is the time required to reach 90% of the final value, whereas *settling time* is the time required to get within some specified amount of the final value and stay there. *Overshoot* and *ringing* are self-explanatory terms for some undesirable properties of filters.

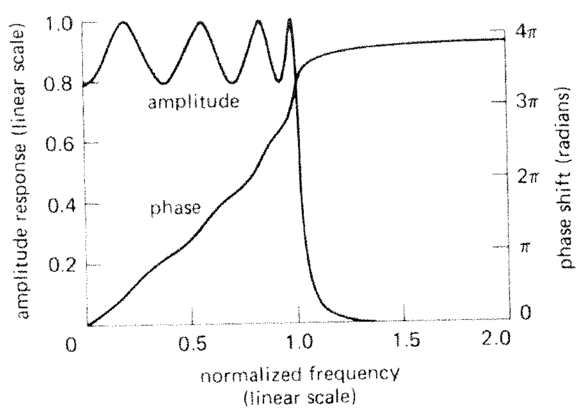

Figure 5.8. Phase and amplitude response for an 8-pole Chebyshev low-pass filter (2dB passband ripple).

5.05 Filter types

Suppose you want a low-pass filter with flat passband and sharp transition to the stopband. The ultimate rate of falloff, well into the stopband, will always be 6ndB/octave, where n is the number of "poles." You need one capacitor (or inductor) for each pole, so the required ultimate rate of falloff of filter response determines, roughly, the complexity of the filter.

Now, assume that you have decided to use a 6-pole low-pass filter. You are guaranteed an ultimate rolloff of 36dB/octave at high frequencies. It turns out

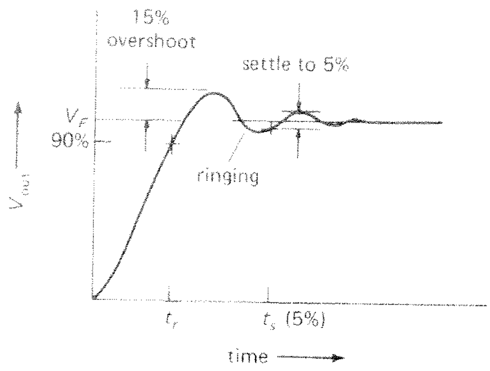

Figure 5.9

that the filter design can now be optimized for maximum flatness of passband response, at the expense of a slow transition from passband to stopband. Alternatively, by allowing some ripple in the passband characteristic, the transition from passband to stopband can be steepened considerably. A third criterion that may be important is the ability of the filter to pass signals within the passband without distortion of their waveforms caused by phase shifts. You may also care about rise time, overshoot, and settling time.

There are filter designs available to optimize each of these characteristics, or combinations of them. In fact, rational filter selection will not be carried out as just described; rather, it normally begins with a set of requirements on passband flatness, attenuation at some frequency outside the passband, and whatever else matters. You will then choose the best design for the job, using the number of poles necessary to meet the requirements. In the next few sections we will introduce the three popular favorites, the Butterworth filter (maximally flat passband), the Chebyshev filter (steepest transition from passband to stopband), and the Bessel filter (maximally flat time delay). Each of these filter responses can be produced with a variety of different filter circuits, some of which we will discuss later. They are all available in low-pass, high-pass, and bandpass versions.

Butterworth and Chebyshev filters

The Butterworth filter produces the flattest passband response, at the expense of steepness in the transition region from passband to stopband. As you will see later, it also has poor phase characteristics. The amplitude response is given by

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{1}{[1 + (f/f_c)^{2n}]^{\frac{1}{2}}}$$

where n is the order of the filter (number of poles). Increasing the number of poles flattens the passband response and steepens the stopband falloff, as shown in Figure 5.10.

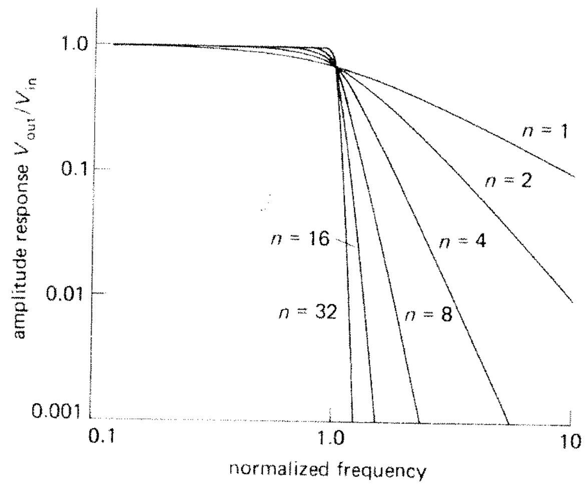

Figure 5.10. Normalized low-pass Butterworth-filter response curves. Note the improved attenuation characteristics for the higher-order filters.

The Butterworth filter trades off everything else for maximum flatness of response. It starts out extremely flat at zero frequency and bends over near the cutoff frequency f_c (f_c is usually the -3dB point).

In most applications, all that really matters is that the wiggles in the passband response be kept less than some amount, say 1dB. The Chebyshev filter responds to this reality by allowing some ripples throughout the passband, with greatly improved sharpness of the knee. A Chebyshev filter is specified in terms of its number of poles and passband ripple. By allowing greater passband ripple, you get a sharper knee. The amplitude is given by

$$\frac{V_{out}}{V_{in}} = \frac{1}{[1+\epsilon^2 C_n^2(f/f_c)]^{\frac{1}{2}}}$$

where C_n is the Chebyshev polynomial of the first kind of degree n, and ϵ is a constant that sets the passband ripple. Like the Butterworth, the Chebyshev has phase characteristics that are less than ideal.

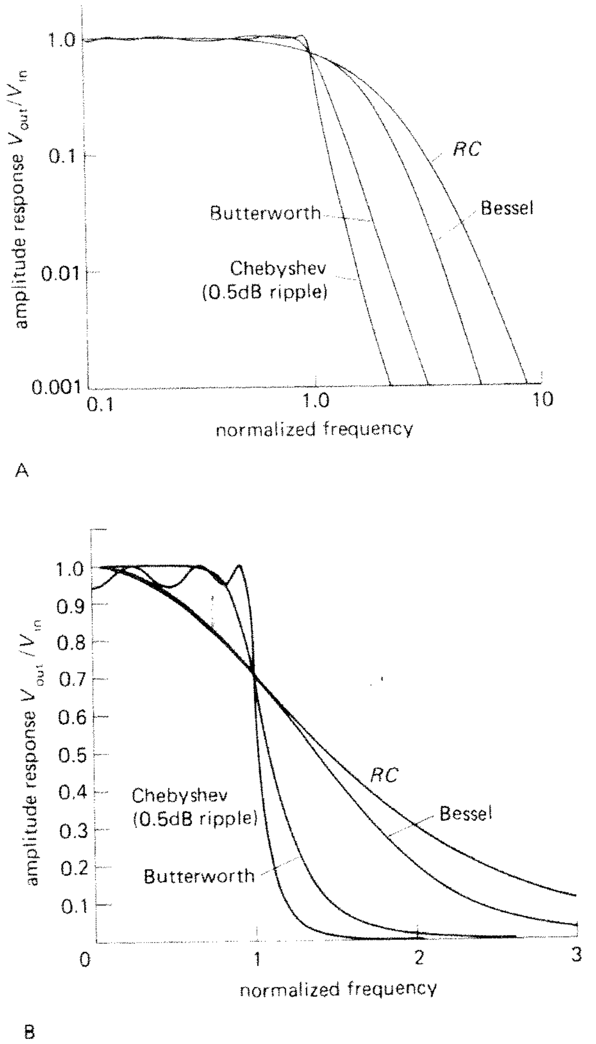

Figure 5.11. Comparison of some common 6-pole low-pass filters. The same filters are plotted on both linear and logarithmic scales.

Figure 5.11 presents graphs comparing the responses of Chebyshev and Butterworth 6-pole low-pass filters. As you can see, they're both tremendous improvements over a 6-pole RC filter.

Actually, the Butterworth, with its maximally flat passband, is not as attractive as it might appear, since you are always accepting some variation in passband response anyway (with the Butterworth it is a gradual rolloff near f_c, whereas with the Chebyshev it is a set of ripples spread throughout the passband). Furthermore, active filters constructed with components of finite tolerance will deviate from the predicted response, which means that a real Butterworth filter will exhibit some passband ripple anyway. The graph in Figure 5.12 illustrates the effects of worst-case variations in resistor and capacitor values on filter response.

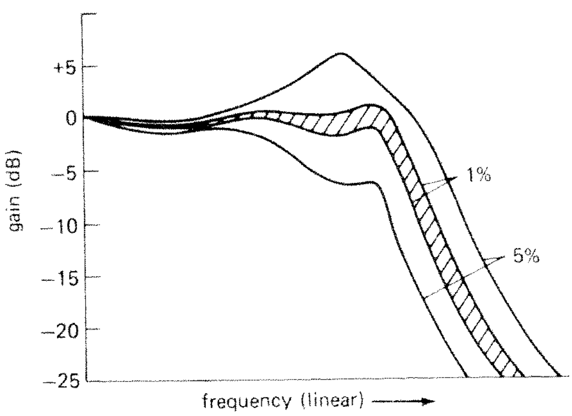

Figure 5.12. The effect of component tolerance on active filter performance.

Viewed in this light, the Chebyshev is a very rational filter design. It is sometimes called an equiripple filter: It manages to improve the situation in the transition region by spreading equal-size ripples throughout the passband, the number of ripples increasing with the order of the filter. Even with rather small ripples (as little as 0.1dB) the Chebyshev filter offers considerably improved sharpness of the knee

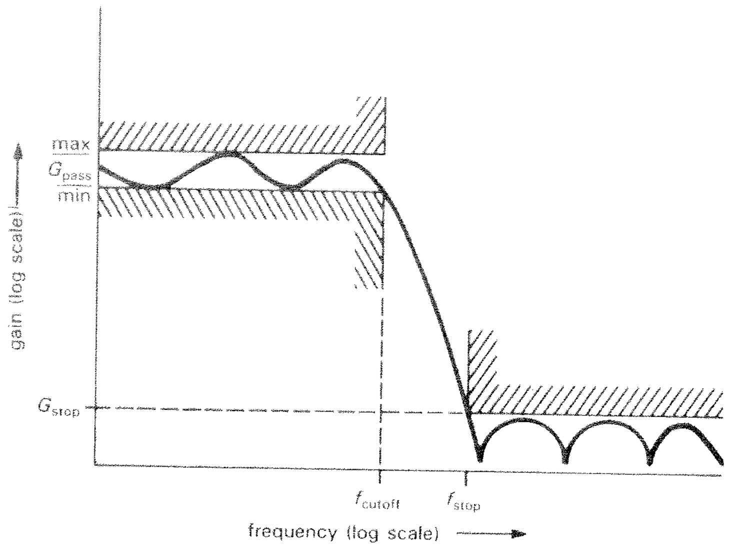

Figure 5.13. Specifying filter frequency response parameters.

as compared with the Butterworth. To make the improvement quantitative, suppose that you need a filter with flatness to 0.1dB within the passband and 20dB attenuation at a frequency 25% beyond the top of the passband. By actual calculation, that will require a 19-pole Butterworth, but only an 8-pole Chebyshev.

The idea of accepting some passband ripple in exchange for improved steepness in the transition region, as in the equiripple Chebyshev filter, is carried to its logical limit in the so-called elliptic (or Cauer) filter by trading ripple in both passband and stopband for an even steeper transition region than that of the Chebyshev filter. With computer-aided design techniques, the design of elliptic filters is as straightforward as for the classic Butterworth and Chebyshev filters.

Figure 5.13 shows how you specify filter frequency response graphically. In this case (a low-pass filter) you indicate the allowable range of filter gain (i.e., the ripple) in the passband, the minimum frequency at which the response leaves the passband, the maximum frequency at which the response enters the stopband, and the minimum attenuation in the stopband.

Bessel filter

As we hinted earlier, the amplitude response of a filter does not tell the whole story. A filter characterized by a flat amplitude response may have large phase shifts. The result is that a signal in the passband will suffer distortion of its waveform. In situations where the shape of the waveform is paramount, a linear-phase filter (or constant-time-delay filter) is desirable. A filter whose phase shift varies linearly with frequency is equivalent to a constant time delay for signals within the passband, i.e., the waveform is not distorted. The Bessel filter (also called the Thomson filter) had maximally flat time delay within its passband, in analogy with the Butterworth, which has maximally flat amplitude response. To see the kind of improvement in time-domain performance you get with the Bessel filter, look at Figure 5.14 for a comparison of time delay versus normalized frequency for 6-pole Bessel and Butterworth low-pass filters. The poor time-delay performance of the Butterworth gives rise to effects such as overshoot when driven with pulse signals. On the other hand, the price you pay for the Bessel's constancy of time delay is an amplitude response with even less steepness than that of the Butterworth in the transition region between passband and stopband.

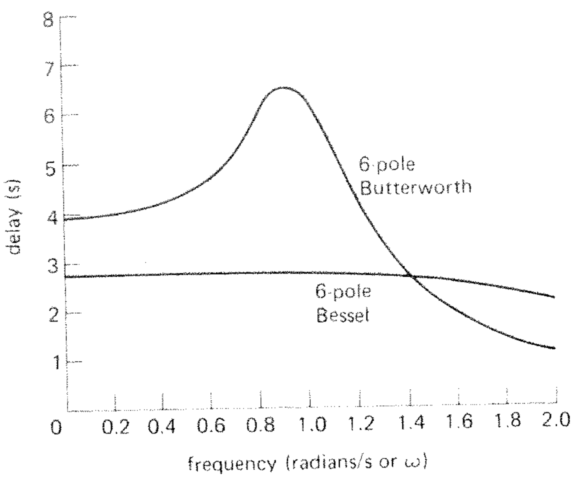

Figure 5.14. Comparison of time delays for 6-pole Bessel and Butterworth low-pass filters. The excellent time-domain performance of the Bessel filter minimizes waveform distortion.

There are numerous filter designs that attempt to improve on the Bessel's good time-domain performance by compromising some of the constancy of time delay for improved rise time and amplitude-versus-frequency characteristics. The Gaussian filter has phase characteristics nearly as good as those of the Bessel, with improved step response. In another class there are interesting filters that allow uniform ripples in the passband time delay (in analogy with the Chebyshev's ripples in its amplitude response) and yield approximately constant time delays even for signals well into the stopband. Another approach to the problem of getting filters with uniform time delays is to use all-pass filters, also known as delay equalizers. These have constant amplitude response with frequency, with a phase shift that can be tailored to individual requirements. Thus, they can be used to improve the time-delay constancy of any filter, including Butterworth and Chebyshev types.

Filter comparison

In spite of the preceding comments about the Bessel filter's transient response, it still has vastly superior properties in the time domain, as compared with the Butterworth and Chebyshev. The Chebyshev, with its highly desirable amplitude-versus-frequency characteristics, actually has the poorest time-domain performance of the three. The Butterworth is in between in both frequency and time-domain properties. Table 5.1 and Figure 5.15 give more information about time-domain performance for these three kinds of filters to complement the frequency-domain graphs presented earlier. They make it clear that the Bessel is a very desirable filter where performance in the time domain is important.

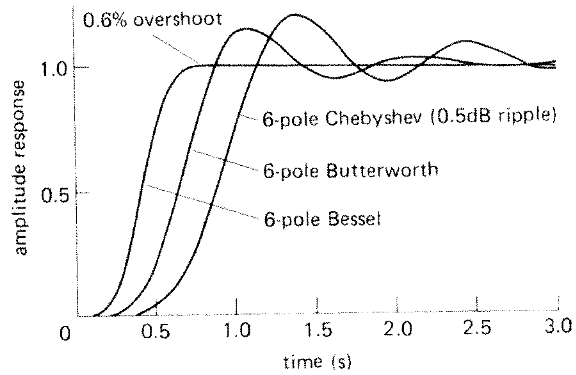

Figure 5.15. Step-response comparison for 6-pole low-pass filters normalized for 3dB attenuation at 1Hz.

ACTIVE FILTER CIRCUITS

A lot of ingenuity has been used in inventing clever active circuits, each of which can be used to generate response functions such as the Butterworth, Chebyshev, etc. You might wonder why the world needs more than one active filter circuit. The reason is that various circuit realizations excel in one or another desirable property, so there is no all-around best circuit.

Some of the features to look for in active filters are (a) small numbers of parts, both

TABLE 5.1. TIME-DOMAIN PERFORMANCE COMPARISON FOR LOW-PASS FILTERS[a]

Type	f_{3dB} (Hz)	Poles	Step rise time (0 to 90%) (s)	Overshoot (%)	Settling time to 1% (s)	Settling time to 0.1% (s)	Stopband attenuation $f = 2f_c$ (dB)	Stopband attenuation $f = 10f_c$ (dB)
Bessel	1.0	2	0.4	0.4	0.6	1.1	10	36
(-3.0dB at	1.0	4	0.5	0.8	0.7	1.2	13	66
f_c = 1.0Hz)	1.0	6	0.6	0.6	0.7	1.2	14	92
	1.0	8	0.7	0.3	0.8	1.2	14	114
Butterworth	1.0	2	0.4	4	0.8	1.7	12	40
(-3.0dB at	1.0	4	0.6	11	1.0	2.8	24	80
f_c = 1.0Hz)	1.0	6	0.9	14	1.3	3.9	36	120
	1.0	8	1.1	16	1.6	5.1	48	160
Chebyshev	1.39	2	0.4	11	1.1	1.6	8	37
0.5dB ripple	1.09	4	0.7	18	3.0	5.4	31	89
(-0.5dB at	1.04	6	1.1	21	5.9	10.4	54	141
f_c = 1.0Hz)	1.02	8	1.4	23	8.4	16.4	76	193
Chebyshev	1.07	2	0.4	21	1.6	2.7	15	44
2.0dB ripple	1.02	4	0.7	28	4.8	8.4	37	96
(-2.0dB at	1.01	6	1.1	32	8.2	16.3	60	148
f_c = 1.0Hz)	1.01	8	1.4	34	11.6	24.8	83	200

(a) a design procedure for these filters is presented in Section 5.07.

active and passive, (b) ease of adjustability, (c) small spread of parts values, especially the capacitor values, (d) undemanding use of the op-amp, especially requirements on slew rate, bandwidth, and output impedance, (e) the ability to make high-Q filters, and (f) sensitivity of filter characteristics to component values and op-amp gain (in particular, the gain-bandwidth product, f_T). In many ways the last feature is one of the most important. A filter that requires parts of high precision is difficult to adjust, and it will drift as the components age; in addition, there is the nuisance that it requires components of good initial accuracy. The VCVS circuit probably owes most of its popularity to its simplicity and its low parts count, but it suffers from high sensitivity to component variations. By comparison, recent interest in more complicated filter realizations is motivated by the benefits of insensitivity of filter properties to small component variability.

In this section we will present several circuits for low-pass, high-pass, and bandpass active filters. We will begin with the popular VCVS, or controlled-source type, then show the state-variable designs available as integrated circuits from several manufacturers, and finally mention the twin-T sharp rejection filter and some interesting new directions in switched-capacitor realizations.

5.06 VCVS circuits

The voltage-controlled voltage-source (VCVS) filter, also known simply as a controlled-source filter, is a variation of the Sallen-and-Key circuit shown earlier. It replaces the unity-gain follower with a noninverting amplifier of gain greater than 1. Figure 5.16 shows the circuits for low-pass, high-pass, and bandpass realizations. The resistors at the outputs of the op-amps create a noninverting voltage amplifier

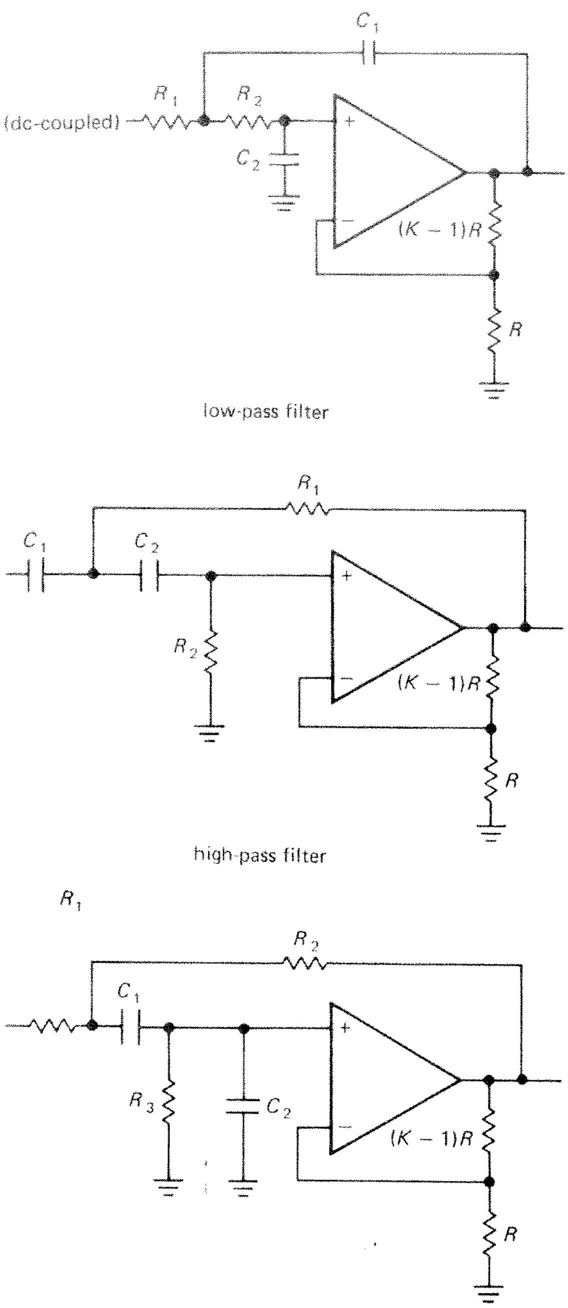

Figure 5.16. VCVS active filter circuits.

of voltage gain K, with the remaining Rs and Cs contributing the frequency response properties for the filter. These are 2-pole filters, and they can be Butterworth, Bessel, etc., by suitable choice of component values, as we will show later. Any number of VCVS 2-pole sections may be cascaded to generate higher-order filters. When that is done, the individual filter sections are, in general, not identical. In fact, each section represents a quadratic polynomial factor of the nth-order polynomial describing the overall filter.

There are design equations and tables in most standard filter handbooks for all the standard filter responses, usually including separate tables for each of a number of ripple amplitudes for Chebyshev filters. In the next section we will present an easy-to-use design table for VCVS filters of Butterworth, Bessel, and Chebyshev responses (0.5dB and 2dB passband ripple for Chebyshev filters) for use as low-pass or high-pass filters. Bandpass and band-reject filters can be easily made from combinations of these.

5.07 VCVS filter design using our simplified table

To use Table 5.2, begin by deciding which filter response you need. As we mentioned earlier, the Butterworth may be attractive if maximum flatness of passband is desired, the Chebyshev gives the fastest rolloff from passband to stopband (at the

TABLE 5.2. VCVS LOW-PASS FILTERS

Poles	Butter-worth	Bessel		Chebyshev (0.5dB)		Chebyshev (2.0dB)	
	K	f_n	K	f_n	K	f_n	K
2	1.586	1.272	1.268	1.231	1.842	0.907	2.114
4	1.152	1.432	1.084	0.597	1.582	0.471	1.924
	2.235	1.606	1.759	1.031	2.660	0.964	2.782
6	1.068	1.607	1.040	0.396	1.537	0.316	1.891
	1.586	1.692	1.364	0.768	2.448	0.730	2.648
	2.483	1.908	2.023	1.011	2.846	0.983	2.904
8	1.038	1.781	1.024	0.297	1.522	0.238	1.879
	1.337	1.835	1.213	0.599	2.379	0.572	2.605
	1.889	1.956	1.593	0.861	2.711	0.842	2.821
	2.610	2.192	2.184	1.006	2.913	0.990	2.946

expense of some ripple in the passband), and the Bessel provides the best phase characteristics, i.e., constant signal delay in the passband, with correspondingly good step response. The frequency responses for all types are shown in the accompanying graphs (Fig. 5.17).

To construct an n-pole filter (n is an even number), you will need to cascade $n/2$ VCVS sections. Only even-order filters are shown, since an odd-order filter requires as many op-amps as the next higher-order filter. Within each section, $R_1 = R_2 = R$, and $C_1 = C_2 = C$. As is usual in op-amp circuits, R will typically be chosen in the range 10k to 100k. (It is best to avoid small resistor values, because the rising open-loop output impedance of the op-amp at high frequencies adds to the resistor values and upsets calculations.) Then all you need to do is set the gain, K, of each stage according to the table entries. For an n-pole filter there are $n/2$ entries, one for each section.

Butterworth low-pass filters

If the filter is a Butterworth, all sections have the same values of R and C, given simply by $RC = 1/2\pi f_c$, where f_c is the desired −3dB frequency of the entire filter. To make a 6-pole low-pass Butterworth filter, for example, you cascade three of the low-pass sections shown previously, with gains of 1.07, 1.59, and 2.48 (preferably in that order, to avoid dynamic range problems), and with identical Rs and Cs to set the 3dB point. The telescope drive circuit in Section 8.31 shows such an example, with f_c = 88.4Hz (R = 180k, $C = 0.01\mu$F).

Bessel and Chebyshev low-pass filters

To make a Bessel or Chebyshev filter with the VCVS, the situation is only slightly more complicated. Again we cascade several 2-pole VCVS filters, with prescribed gains for each section. Within each section we again use $R_1 = R_2 = R$, and $C_1 = C_2 = C$. However, unlike the situation with the Butterworth, the RC products for the different sections are different and must be scaled by the normalizing factor f_n (given for each section in Table 5.2) according to $RC = 1/2\pi f_n f_c$. Here f_c is again the −3dB point for the Bessel filter, whereas for the Chebyshev filter it defines the end of the passband, i.e., it is the frequency at which the amplitude response falls out of the ripple band on its way into the stopband. For example, the response of a Chebyshev low-pass filter with 0.5dB ripple and f_c = 100Hz will be flat within +0dB to −0.5dB from dc to 100Hz, with 0.5dB attenuation at 100Hz and a rapid falloff for frequencies greater than 100Hz. Values are given for Chebyshev filters with 0.5dB and 2.0dB passband ripple; the latter have a somewhat steeper transition into the stopband (Fig. 5.17).

High-pass filters

To make a high-pass filter, use the high-pass configuration shown previously, i.e., with the Rs and Cs interchanged. For Butterworth filters, everything else remains unchanged (use the same values for R, C, and K). For the Bessel and Chebyshev filters, the K values remain the same, but the normalizing factors f_n must be inverted, i.e., for each section the new f_n equals $1/(f_n$ listed in Table 5.2).

A bandpass filter can be made by cascading overlapping low-pass and high-pass filters. A band-reject filter can be made by summing the outputs of nonoverlapping low-pass and high-pass filters. However, such cascaded filters won't work well for high-Q filters (extremely sharp bandpass filters) because there is great sensitivity to the component values in the individual (uncoupled) filter sections. In such cases a high-Q single-stage bandpass circuit (e.g., the VCVS bandpass circuit

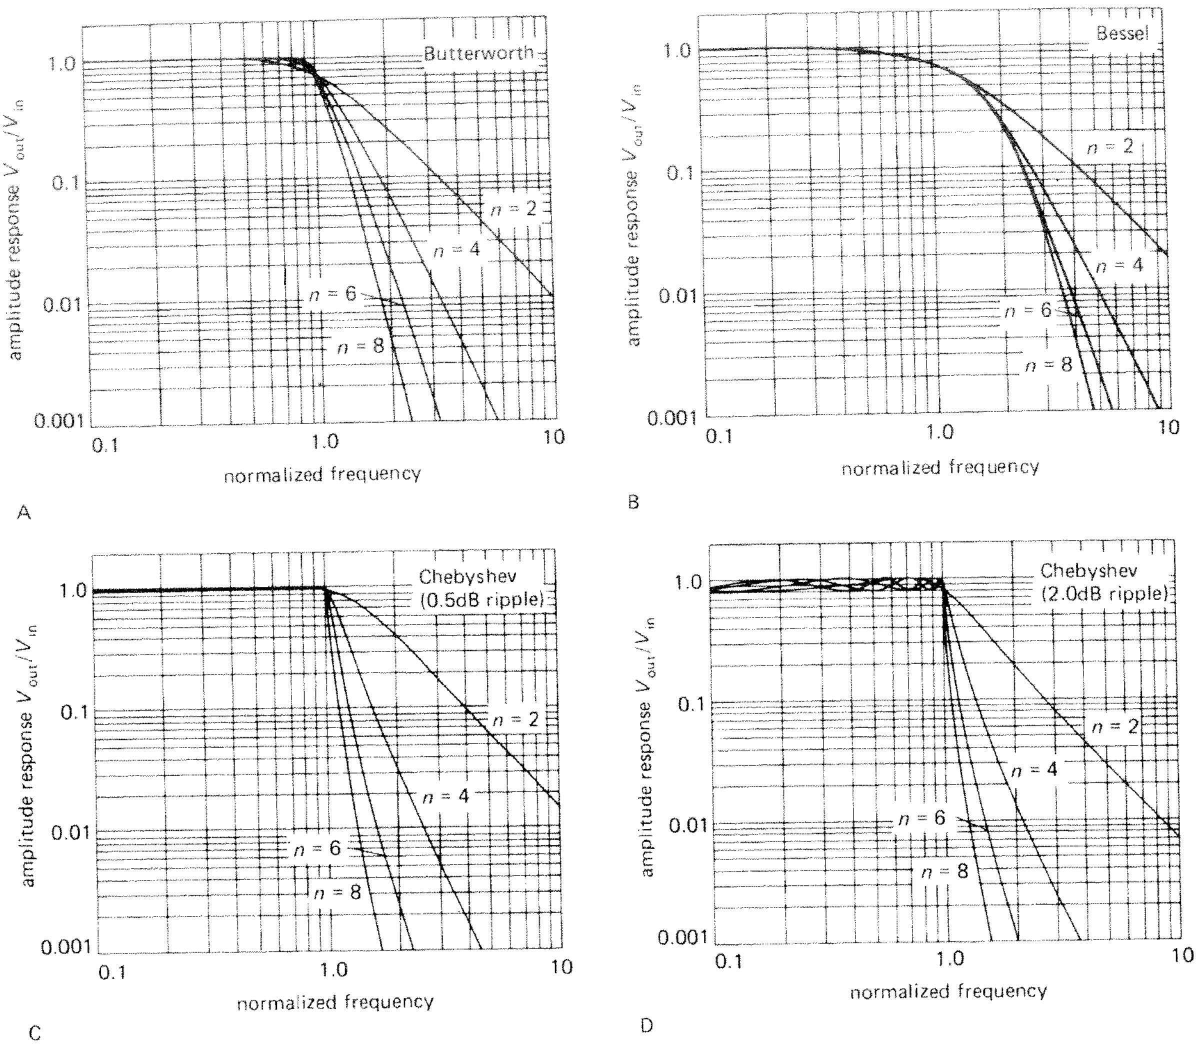

Figure 5.17. Normalized frequency response graphs for the 2-, 4-, 6-, and 8-pole filters in Table 5.2. The Butterworth and Bessel filters are normalized to 3dB attenuation at unit frequency, whereas the Chebyshev filters are normalized to 0.5dB and 2dB attenuations.

illustrated previously, or the state-variable and biquad filters in the next section) should be used instead. Even a single-stage 2-pole filter can produce a response with an extremely sharp peak. Information on such filter design is available in the standard references.

VCVS filters minimize the number of components needed (2 poles/op-amp) and offer the additional advantages of noninverting gain, low output impedance, small spread of component values, easy adjustability of gain, and the ability to operate at high gain or high Q. They suffer from high sensitivity to component values and amplifier gain, and they don't lend themselves well to applications where a tunable filter of stable characteristics is needed.

EXERCISE 5.3

Design a 6-pole Chebyshev low-pass VCVS filter with a 0.5dB passband ripple and 100Hz cutoff frequency f_c. What is the attenuation at $1.5f_c$?

5.08 State-variable filters

The 2-pole filter shown in Figure 5.18 is far more complex than the VCVS circuits,

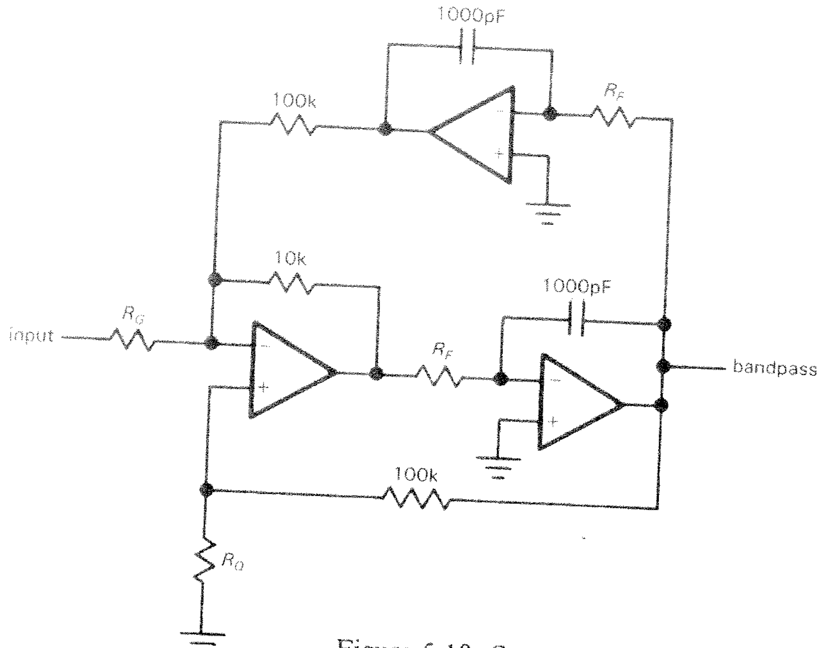

Figure 5.18. State-variable active filter.

but it is popular because of its improved stability and ease of adjustment. It is called a state-variable filter and is available as an IC from National (the AF100 and AF150), Burr-Brown (the UAF series), and others. Because it is a manufactured module, all components except R_G, R_Q, and the two R_Fs are built in. Among its nice properties is the availability of high-pass, low-pass, and bandpass outputs from the same circuit; also, its frequency can be tuned while maintaining constant Q (or, alternatively, constant bandwidth) in the bandpass characteristic. As with the VCVS realizations, multiple stages can be cascaded to generate higher-order filters.

Extensive design formulas and tables are provided by the manufacturers for the use of these convenient ICs. They show how to choose the external resistor values to make Butterworth, Bessel, and Chebyshev filters for a wide range of filter orders, for low-pass, high-pass, bandpass, and band-reject responses. Among the nice features of these hybrid ICs is integration of the capacitors into the module, so that only external resistors need be added.

Bandpass filters

The state-variable circuit, in spite of its large number of components, is a good choice for sharp (high-Q) bandpass filters. It has low component sensitivities, does not make great demands on op-amp bandwidth, and is easy to tune. For example, in the circuit of Figure 5.18, used as a bandpass filter, the two resistors R_F set the center frequency, while R_Q and R_G together determine the Q and band-center gain:

$$R_F = 5.03 \times 10^7/f_0 \text{ ohms}$$

$$R_Q = 10^5/(3.48Q + G - 1) \text{ ohms}$$

$$R_G = 3.16 \times 10^4 Q/G \text{ ohms}$$

So you could make a tunable-frequency, constant-Q filter by using a 2-section variable resistor (pot) for R_F. Alternatively, you could make R_Q adjustable, producing a fixed-frequency, variable-Q (and, unfortunately, variable-gain) filter.

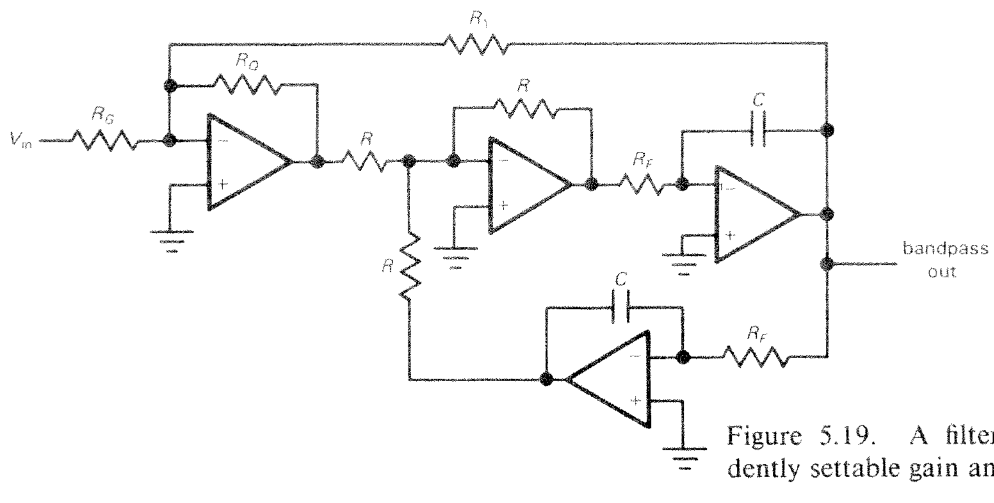

Figure 5.19. A filter with independently settable gain and Q.

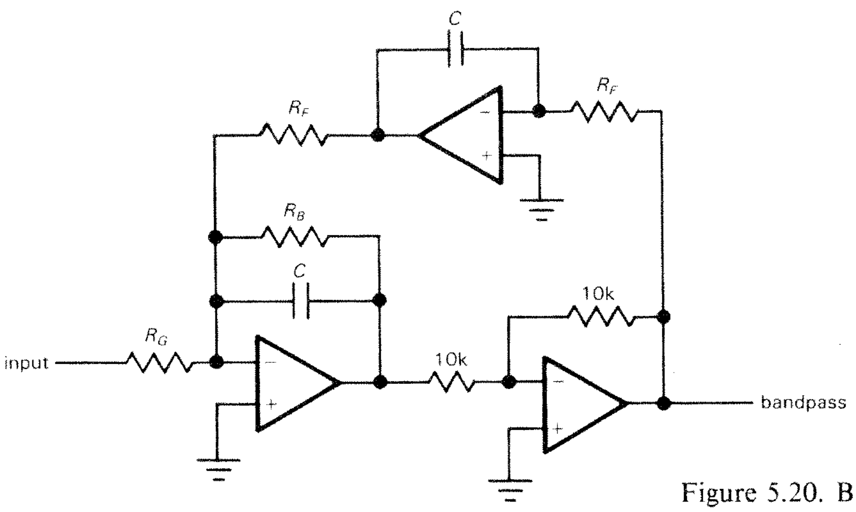

Figure 5.20. Biquad active filter.

EXERCISE 5.4

Calculate resistor values in Figure 5.18 to make a bandpass filter with $f_0 = 1\text{kHz}$, $Q = 50$, and $G = 10$.

Figure 5.19 shows a useful variant of the state-variable bandpass filter. The bad news is that it uses four op-amps; the good news is that you can adjust the bandwidth (i.e., Q) without affecting the midband gain. In fact, both Q and gain are set with a single resistor each. Q, gain, and center frequency are completely independent and are given by these simple equations:

$f_0 = 1/2\pi R_F C$
$Q = R_1/R_Q$
$G = R_1/R_G$
$R \approx 10\text{k}$ (noncritical, matched)

Biquad filter. A close relative of the state variable filter is the so-called biquad filter, shown in Figure 5.20. This circuit also uses three op-amps and can be constructed from the state-variable ICs mentioned earlier. It has the interesting property that you can tune its frequency (via R_F) while maintaining constant *bandwidth* (rather than constant Q). Here are the design equations:

$f_0 = 1/2\pi R_F C$
$\text{BW} = 1/2\pi R_B C$
$G = R_B/R_G$

The Q is given by f_0/BW and equals R_B/R_F. As the center frequency is varied (via R_F), the Q varies proportionately, keeping the bandwidth Qf_0 constant.

When you design a biquad filter from scratch (rather than with an active filter IC that already contains most of the parts), the general procedure goes something like this:

1. Choose an op-amp whose bandwidth f_T is at least 10 to 20 times Gf_0.

2. Pick a round-number capacitor value in the vicinity of

$$C = 10/f_0\ \mu F$$

3. Use the desired center frequency to calculate the corresponding R_F from the first equation given earlier.

4. Use the desired bandwidth to calculate R_B from the second equation given earlier.

5. Use the desired band-center gain to calculate R_G from the third equation given earlier.

You may have to adjust the capacitor value if the resistor values become awkwardly large or small. For instance, in a high-Q filter you may need to increase C somewhat to keep R_B from becoming too large (or you can use the T-network trick described in Section 4.19). Note that R_F, R_B, and R_G each act as op-amp loads, and should not become less than, say, 5k. When juggling component values, you may find it easier to satisfy requirement 1 by decreasing integrator gain (increase R_F) and simultaneously increasing the inverter-stage gain (increase the 10k feedback resistor).

As an example, suppose we want to make a filter with the same characteristics as in the last exercise. We would begin by provisionally choosing $C = 0.01\mu F$. Then we find $R_F = 15.9k$ ($f_0 = 1kHz$) and $R_B = 796k$ ($Q = 50$; BW=20Hz). Finally, $R_G = 79.6k$ ($G = 10$).

EXERCISE 5.5

Design a biquad bandpass filter with $f_0 = 60Hz$, BW=1Hz, and $G = 100$.

Higher order bandpass filters

As with our earlier low-pass and high-pass filters, it is possible to build higher order bandpass filters with approximately flat bandpass and steep transition to the stopband.

You do this by cascading several lower-order bandpass filters, the combination tailored to realize the desired filter type (Butterworth, Chebyshev, or whatever). As before, the Butterworth is "maximally flat," whereas the Chebyshev sacrifices passband flatness for steepness of skirts. Both the VCVS and state-variable/biquad bandpass filters just considered are second order (two pole). As you increase the filter sharpness by adding sections, you generally degrade the transient response and phase characteristics. The "bandwidth" of a bandpass filter is defined as the width between −3dB points, except for equiripple filters, for which it is the width between frequencies at which the response falls out of the passband ripple channel.

You can find tables and design procedures for constructing complex filters in standard books on active filters, or in the data sheets for active filter ICs. There are also some very nice filter design programs that run on inexpensive workstations (IBM PC, Macintosh).

□ 5.09 Twin-T notch filters

The passive RC network shown in Figure 5.21 has infinite attenuation at a frequency $f_c = 1/2\pi RC$. Infinite attenuation is

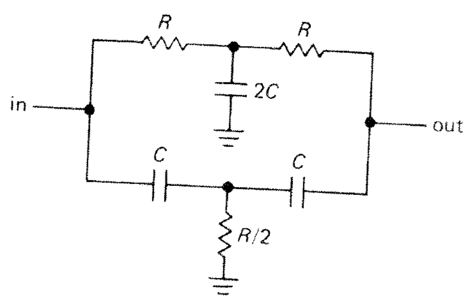

Figure 5.21. Passive twin-T notch filter.

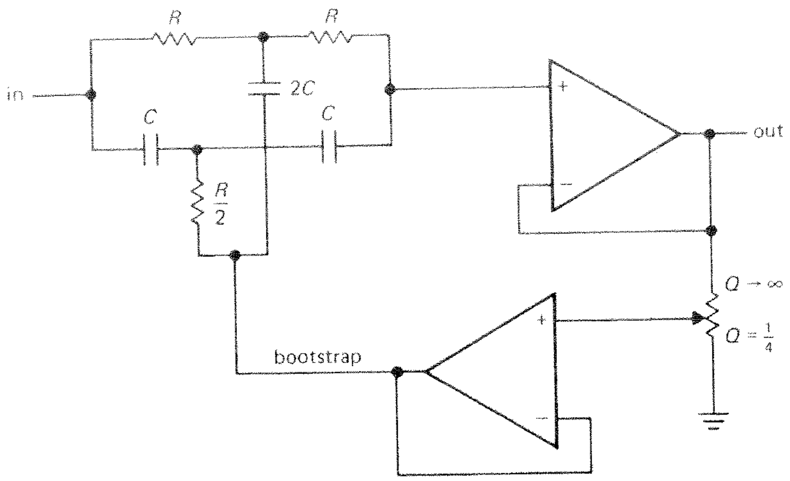

Figure 5.22. Bootstrapped twin-T.

uncharacteristic of *RC* filters in general; this one works by effectively adding two signals that have been shifted 180° out of phase at the cutoff frequency. It requires good matching of components in order to obtain a good null at f_c. It is called a twin-T, and it can be used to remove an interfering signal, such as 60Hz powerline pickup. The problem is that it has the same "soft" cutoff characteristics as all passive *RC* networks, except, of course, near f_c, where its response drops like a rock. For example, a twin-T driven by a perfect voltage source is down 10dB at twice (or half) the notch frequency and 3dB at four times (or one-fourth) the notch frequency. One trick to improve its notch characteristic is to "activate" it in the manner of a Sallen-and-Key filter (Fig. 5.22). This technique looks good in principle, but it is generally disappointing in practice, owing to the impossibility of maintaining a good filter null. As the filter notch becomes sharper (more gain in the bootstrap), its null becomes less deep.

Twin-T filters are available as prefab modules, going from 1Hz to 50kHz, with notch depths of about 60dB (with some deterioration at high and low temperatures). They are easy to make from components, but resistors and capacitors of good stability and low temperature coefficient should be used to get a deep and stable notch. One of the components should be made trimmable.

The twin-T filter works fine as a fixed-frequency notch, but it is a horror to make tunable, since three resistors must be simultaneously adjusted while maintaining constant ratio. However, the remarkably simple *RC* circuit of Figure 5.23A, which behaves just like the twin-T, can be adjusted over a significant range of frequency (at least two octaves) with a single potentiometer. Like the twin-T (and most active filters) it requires some matching of components; in this case the three capacitors must be identical, and the fixed resistor must be exactly six times the bottom (adjustable) resistor. The notch frequency is then given by

$$f_{\text{notch}} = 1/2\pi C\sqrt{3R_1R_2}$$

Figure 5.23B shows an implementation that is tunable from 25Hz to 100Hz. The 50k trimmer is adjusted (once) for maximum depth of notch.

As with the passive twin-T, this filter (known as a *bridged differentiator*) has a gently sloping attenuation away from the notch and infinite attenuation (assuming perfect matching of component values) at the notch frequency. It, too, can be "activated," by bootstrapping the wiper of the pot with a voltage gain somewhat less than unity (as in Fig. 5.22). Increasing

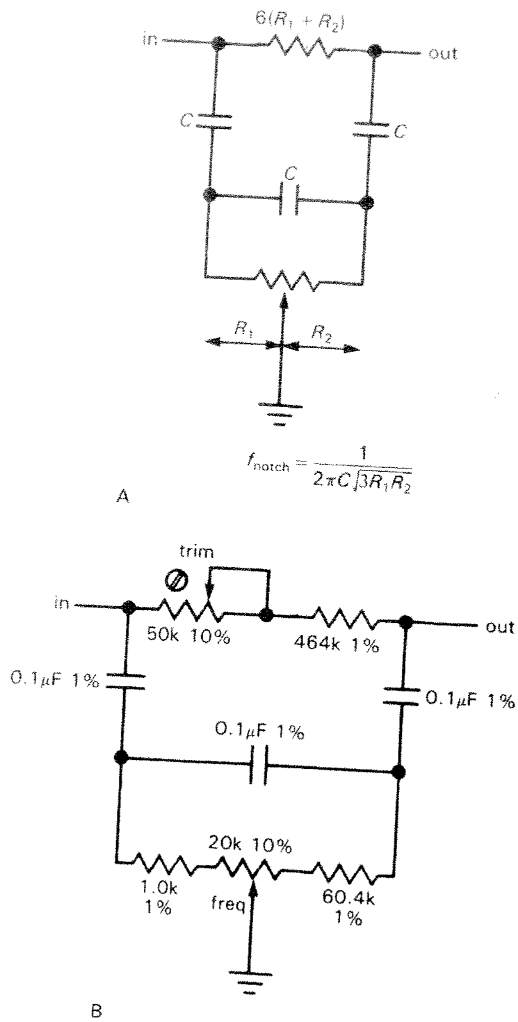

Figure 5.23. Bridged differentiator tunable-notch filter. The implementation in B tunes from 25Hz to 100Hz.

the bootstrap gain toward unity narrows the notch, but also leads to an undesirable response peak on the high frequency side of the notch, along with a reduction in ultimate attenuation.

☐ 5.10 Gyrator filter realizations

An interesting type of active filter is made with gyrators; basically they are used to substitute for inductors in traditional filter designs. The gyrator circuit shown in Figure 5.24 is popular. Z_4 will ordinarily be a capacitor, with the other impedances

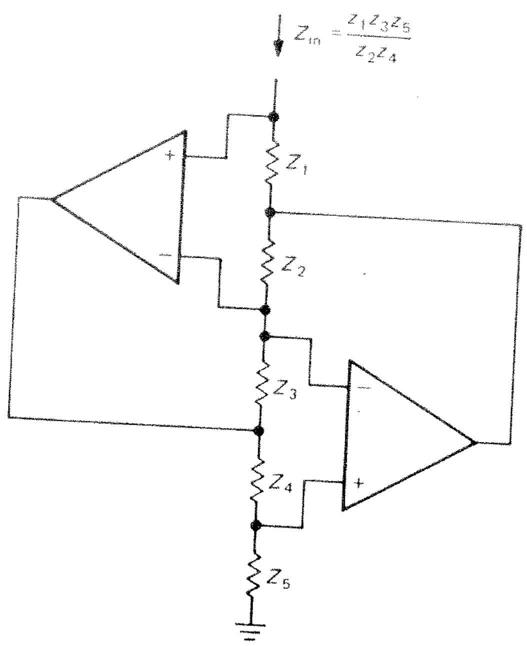

Figure 5.24. Gyrator.

being replaced by resistors, creating an inductor $L = kC$, where $k = R_1R_3R_5/R_2$. It is claimed that these gyrator-substituted filters have the lowest sensitivity to component variations, exactly analogous to their passive *RLC* prototypes.

5.11 Switched capacitor filters

One drawback to these state-variable or biquad filters is the need for accurately matched capacitors. If you build the circuit from op-amps, you've got to get pairs of stable capacitors (not ceramic or electrolytic), perhaps matched as closely as 2% for optimum performance. You also have to make a lot of connections, since the circuits use at least three op-amps and six resistors for each 2-pole section. Alternatively, you can buy a filter IC, letting the manufacturer figure out how to integrate matched 1000pF capacitors into his IC. IC manufacturers have solved those problems, but at a price: The AF100 "Universal Active Filter" IC from National is a hybrid IC and costs about $10 apiece.

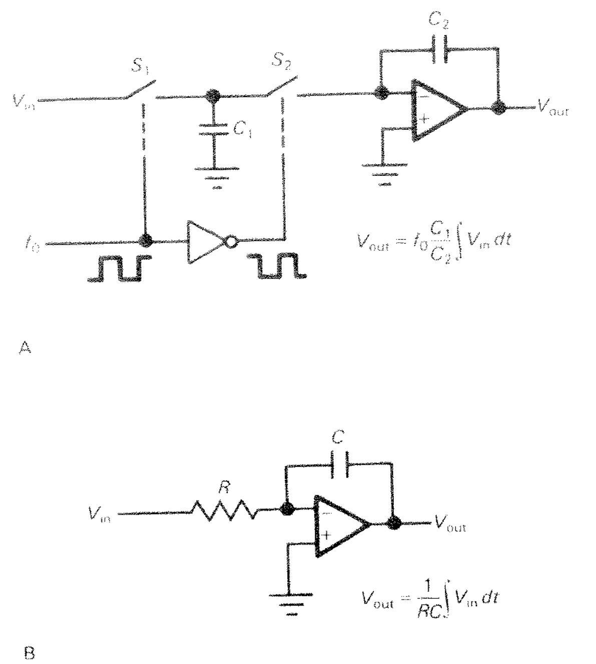

Figure 5.25. A. Switched-capacitor integrator B. conventional integrator.

There's another way to implement the integrators that are needed in the state-variable or biquad filter. The basic idea is to use MOS analog switches, clocked from an externally applied square wave at some high frequency (typically 100 times faster than the analog signals of interest), as shown in Figure 5.25. In the figure, the funny triangular object is a digital *inverter*, which turns the square wave upside down so that the two MOS switches are closed on opposite halves of the square wave. The circuit is easy to analyze: When S_1 is closed, C_1 charges to V_{in}, i.e., holding charge C_1V_{in}; on the alternate half of the cycle, C_1 discharges into the virtual ground, transferring its charge to C_2. The voltage across C_2 therefore changes by an amount $\Delta V = \Delta Q/C_2 = V_{in}C_1/C_2$. Note that the output voltage *change* during each cycle of the fast square wave is proportional to V_{in} (which we assume changes only a small amount during one cycle of square wave), i.e., the circuit is an integrator! It is easy to show that the integrators obey the equations in the figure.

EXERCISE 5.6
Derive the equations in Figure 5.25

There are two important advantages to using switched capacitors instead of conventional integrators. First, as hinted earlier, it can be less expensive to implement on silicon: The integrator gain depends only on the *ratio* of two capacitors, not on their individual values. In general it is easy to make a matched pair of anything on silicon, but very hard to make a similar component (resistor or capacitor) of precise value and high stability. As a result, monolithic switched-capacitor filter ICs are very inexpensive – National's universal switched-capacitor filter (the MF10) costs $2 (compared with $10 for the conventional AF100) and furthermore gives you *two* filters in one package!

The second advantage of switched-capacitor filters is the ability to tune the filter's frequency (e.g., the center frequency of a bandpass filter, or the −3dB point of a low-pass filter) by merely changing the frequency of the square wave ("clock") input. This is because the characteristic frequency of a state-variable or biquad filter is proportional to (and depends only on) the integrator gain.

Switched-capacitor filters are available in both dedicated and "universal" configurations. The former are prewired with on-chip components to form bandpass or low-pass filters, while the latter have various intermediate inputs and outputs brought out so you can connect external components to make anything you want. The price you pay for universality is a larger IC package and the need for external resistors. For example, National's self-contained MF4 Butterworth low-pass filter comes in an 8-pin DIP ($1.30), while their MF5 universal filter comes in a 14-pin DIP ($1.45), requiring 2 or 3 external resistors (depending on which filter configuration you choose). Figure 5.26 shows just how easy it is to use the dedicated type.

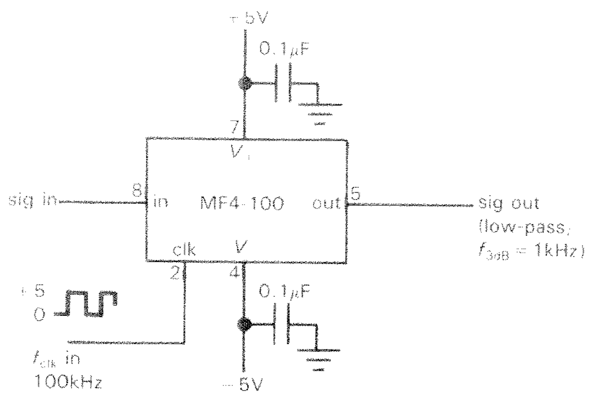

Figure 5.26

Now for the bombshell: Switched-capacitor filters have three annoying characteristics, all related and caused by the presence of the periodic clocking signal. First, there is *clock feedthrough*, the presence of some output signal (typically about 10mV to 25mV) at the clock frequency, independent of the input signal. Usually this doesn't matter, because it is far removed from the signal band of interest. If clock feedthrough is a problem, a simple *RC* filter usually gets rid of it. The second problem is more subtle: If the input signal has any frequency components near the clock frequency, they will be "aliased" down into the passband. To state it precisely, any input signal energy at a frequency that differs from the clock frequency by an amount corresponding to a frequency in the passband will appear (unattenuated!) in the passband. For example, if you use an MF4 as a 1kHz low-pass filter (i.e., set f_{clock} = 100kHz), any input signal energy in the range of 99kHz–101kHz will appear in the output band of dc–1kHz. No filter at the output can remove it! You must make sure the input signal doesn't have energy near the clock frequency. If this isn't naturally the case, you can usually use a simple *RC* filter, since the clock frequency is typically quite far removed from the passband. The third undesirable effect in switched-capacitor filters is a general reduction in signal dynamic range (an increase in the "noise floor") due to incomplete cancellation of MOS switch charge injection (see Section 3.12). Typical filter ICs have dynamic ranges of 80dB–90dB.

Like any linear circuit, switched-capacitor filters (and their op-amp analogs) suffer from amplifier errors such as input offset voltage and $1/f$ low-frequency noise. These can be a problem if, for example, you wish to low-pass filter some low-level signal without introducing errors or fluctuations in its average dc value. A nice solution is provided by the clever folks at Linear Technology, who dreamed up the LTC1062 "DC Accurate Low-Pass Filter" (or the MAX280, with improved offset voltage). Figure 5.27 shows how you use it. The basic idea is to put the filter outside the dc path, letting the low-frequency signal components couple passively to the output; the filter grabs onto the signal line only at higher frequencies, where it rolls off the response by shunting the signal to

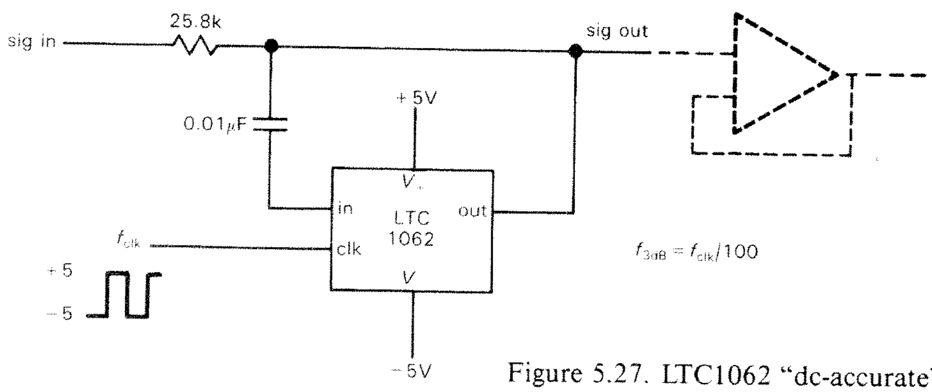

Figure 5.27. LTC1062 "dc-accurate" low-pass filter.

ground. The result is zero dc error, and switched-capacitor-type noise only in the vicinity of the rolloff (Fig. 5.28).

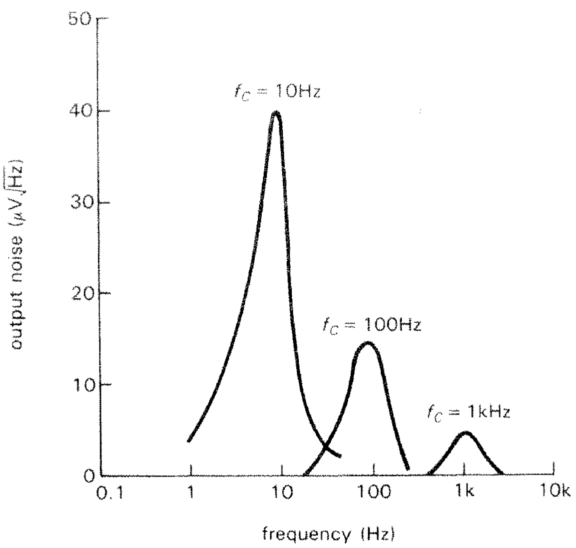

Figure 5.28

Switched-capacitor filter ICs are widely available, from manufacturers such as AMI-Gould, Exar, LTC, National, and EGG-Reticon. Typically you can put the cutoff (or band center) anywhere in the range of dc to a few tens of kilohertz, as set by the clock frequency. The characteristic frequency is a fixed multiple of the clock, usually $50f_{clk}$ or $100f_{clk}$. Most switched-capacitor filter ICs are intended for low-pass, bandpass, or notch (band-stop) use, though a few (e.g., the AMI 3529) are designed as high-pass filters. Note that clock feedthrough and discrete (clock frequency) output waveform quantization effects are particularly bothersome in the latter case, since they're both in-band.

OSCILLATORS

5.12 Introduction to oscillators

Within nearly every electronic instrument it is essential to have an oscillator or waveform generator of some sort. Apart from the obvious case of signal generators, function generators, and pulse generators themselves, a source of regular oscillations is necessary in any cyclical measuring instrument, in any instrument that initiates measurements or processes, and in any instrument whose function involves periodic states or periodic waveforms. That includes just about everything. For example, oscillators or waveform generators are used in digital multimeters, oscilloscopes, radiofrequency receivers, computers, every computer peripheral (tape, disk, printer, alphanumeric terminal), nearly every digital instrument (counters, timers, calculators, and anything with a "multiplexed display"), and a host of other devices too numerous to mention. A device without an oscillator either doesn't do anything or expects to be driven by something else (which probably contains an oscillator). It is not an exaggeration to say that an oscillator of some sort is as essential an ingredient in electronics as a regulated supply of dc power.

Depending on the application, an oscillator may be used simply as a source of regularly spaced pulses (e.g., a "clock" for a digital system), or demands may be made on its stability and accuracy (e.g., the time base for a frequency counter), its adjustability (e.g., the local oscillator in a transmitter or receiver), or its ability to produce accurate waveforms (e.g., the horizontal-sweep ramp generator in an oscilloscope).

In the following sections we will treat briefly the most popular oscillators, from the simple RC relaxation oscillators to the stable quartz-crystal oscillators. Our aim is not to survey everything in exhaustive detail, but simply to make you acquainted with what is available and what sorts of oscillators are suitable in various situations.

5.13 Relaxation oscillators

A very simple kind of oscillator can be made by charging a capacitor through a