J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*

PROGRAMMABLE FILTERS

This article is about programmable electric filter circuits. It is a practical review and involves no difficult mathematics, although it does assume some familiarity with filter terminology. It does not cover digital filters, which tend to be programmable by definition, since they are usually the direct result of running a computer program or at least some kind of number-crunching algorithm.

Programmable filters have characteristics which can be intentionally adjusted to provide useful and desired results, by means of external control inputs which can be of many different forms, both analog and digital.

The circuits described in detail here are programmable in the sense that the resulting response of the filter is a known function of the adjustment process employed. A wider class of filters, which might be termed "variable" rather than "programmable," have response parameters which can be changed by altering something, but perhaps in a way which is an imprecise function of variables under the user's control, or which varies in an unknown way between different units from a production batch; these techniques are not covered in depth here.

The Need for Programmability in a Filter Circuit

In many applications, the response shape, passbands, and stopbands of filter circuits needed in the design can be specified in advance; these frequencies are fixed numbers, the same for every unit which is manufactured and unchanging no matter what the circumstances of use. For such applications, fixed frequency filters are utilized, and design procedures for such circuits are well established.

Applications involving some form of filtering are familiar to everybody; for instance, a radio receiver can be viewed as a form of filter because its job is to pick out a carrier signal on which some wanted information has been impressed by the sender, in the presence of huge amounts of competing unwanted information. A radio receiver which could only receive one station out of the thousands transmitted would be of little use, and so any usable "radio" can in this sense be thought of as a programmable filter.

There are many other applications where filtering is required, in order to discriminate between some wanted signals and some unwanted ones on the basis of their frequencies, but where the characteristics of this filtering need to be changed in response to some changing mode of application or in response to an external circumstance.

The primary form of parameter adjustment covered here will be that of the "cutoff frequency" of the filter; in other words, it is assumed that when the filter is adjusted, its intrinsic response "shape" is not to be altered. For example, if the desired response shape is described by a Chebychev lowpass function, then the flatness and ripple in the passband should not be affected simply by changing the actual value of the passband cutoff frequency. In some cases, programmability of the response shape itself is also required, such as in the programmable audio equalizer described subsequently.

The choice between a fully "programmable" filter , or simply a circuit in which some parameter can be varied somewhat, generally depends on whether it is possible to determine, by examining the result of filtering the signal, whether the filtering was effective. Another way of saying this is that if a form of "closed-loop control" is possible, then it may be possible to get away with an unpredictable adjustment method. If only open-loop operation is possible, then the adjusting technique needs to be correct because an iterative approach to getting the response correct is not possible.

Another frequent feature of the need for programmability, as distinct in this case from "selectability," is that the actual operational parameter required—for instance the cutoff frequency—may not be known until the actual point in time at which it has to be set.

To return again to the radio analogy, a closed-loop technique—listening to the radio and judging whether the desired station is correctly tuned—enables a rough-and-ready adjustment method to be used, one for which the manufacturer may not have made any promises about what the position of the tuning dial actually "means." In contrast, a more modern radio might enable you to "punch in" the frequency of the station with such precision, and to set it with such accuracy, that the tuning is achieved with no correction necessary. This is open-loop operation and requires different (and more accurate) techniques whether the system is a radio or, say, an industrial data acquisition process.

Programmable Elements Available for Filter Parameter Adjustment

We now assume that we are to attempt to control the behavior of a single filter circuit in some way. Naturally, the "brute force" method of achieving a range of filter responses is to design a fixed filter circuit for each of the desired outcome responses and build them all into the equipment, with a big selector switch. This is generally likely to work (assuming no unexpected interactions between the circuits) but is rarely likely to be efficient for component cost, power consumption, or board area consumed and thus it is not usually a serious approach.

The adjustability might be offered to the user in the form of a big front-panel dial, or some form of interface port to attach to a processor or network. Whatever way the interface is achieved, at some point in the circuit programmable components of some sort will be required in order to turn the needs of the user into intentionally adjusted parameters of a filter circuit.

Adjustable or Variable Components. These components generally have no role to play in programmable rather than adjustable filters, but they are useful in certain noncritical circuits. This is because the control relationship may be known only empirically and may offer poor consistency over batches and perhaps unknown stability over time and temperature.

Examples are (1) light-sensitive resistors using cadmium sulfide or a semiconductor junction, programmed using a controllable light source, and (2) rotary resistive potentiometers adjusted with a motor (some form of positional feedback might be used to improve the accuracy and consistency of this approach).

Note that discrete junction field-effect transistors (*FET*s) make excellent switches for the techniques of the next sections, but their analog characteristics—the region between fully on and fully off—are not well-defined and so it is difficult to make repeatable circuits when used as variable resistances with a programmable gate-source voltage. However, careful matching of individual FETs may produce usable results over a range of up to a decade (1,2).

Switched Two-terminal Component Arrays. These comprise arrays of fixed two-terminal components (i.e., resistors, capacitors, or inductors) connected together by controllable switches to provide composite components whose value can be selected from a range of values (but which is constant any time that the selection is not being changed).

This approach is perhaps the most obvious one; it is similar to the "build several complete filters and select between them" method except that it changes less than the entire network of filter components each time you select a new cutoff frequency. How many components you actually need to change when you change

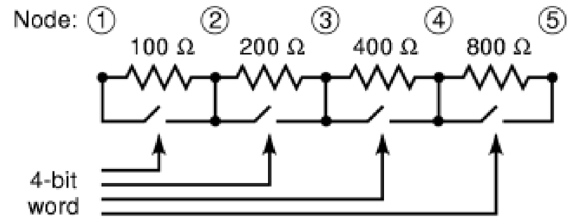

Fig. 1. A 4-bit binary-weighted resistor array.

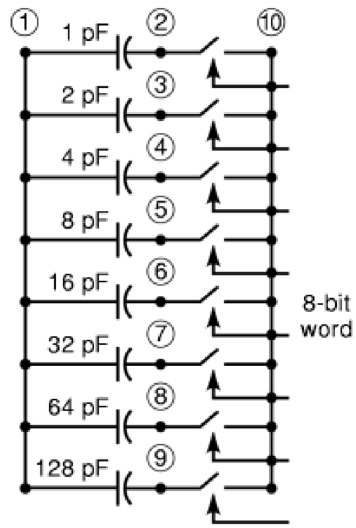

Fig. 2. An 8-bit binary-weighted capacitor array.

the cutoff frequency depends on the filter circuit used; clearly, the fewer the better; this is covered in more detail in the section entitled "Fundamental Restrictions on Programmability."

Binary-Weighted Arrays. A particularly useful technique for constructing programmable components—and it can be capacitors (3) or resistors—is to use a switched array of components with a binary weighting between the values. A couple of examples will illustrate the point.

Firstly, consider four resistors in series, of value 100 Ω, 200 Ω, 400 Ω, and 800 Ω. Place a controllable switch in parallel with each resistor, and construct a truth table showing the realized resistance from node 1 to node 5 for each of the 16 possible 4-bit input control "words" which are used to control the switches. This shows that the circuit generates a resistance whose value is 100 Ω multiplied by the value of this digital word—from 0 Ω to 1500 Ω in increments of 100 Ω, assuming the switches' resistance is negligible (Fig. 1).

Another example is eight capacitors, of value 1 pF, 2 pF, 4 pF, and so on up to 128 pF, each in series with a switch, such that all capacitors share one common terminal (node 1) and all switches share another (node 10) as shown in Fig. 2. A 256-entry truth table could be constructed, though this is unnecessary because it should be fairly clear that this circuit will realize a programmable capacitance which can be varied from 0 pF (all switches open) to 255 pF (all switches closed) in steps of 1 pF. Here we are ignoring any capacitance contribution from the switch elements, which at this level of capacitance can be quite significant.

"Programmable" two-terminal components can be substituted quite freely in most active (and passive) filter circuits for conventional components, as long as the limitations of the switch technology do not affect the circuit performance.

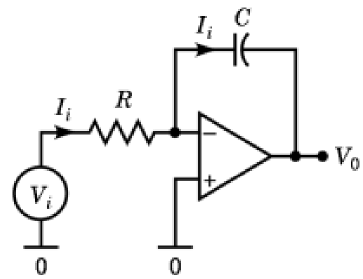

Fig. 3. An inverting integrator.

The Integrator. Before moving on to the other techniques in this section, let us introduce a versatile circuit block with great relevance to active filter design, the integrator. The simplest form, the operational amplifier-based integrator (Fig. 3), uses the "virtual earth" at the inverting input of an op-amp operating with feedback, to produce an output voltage which is the time integral of the current flowing through the capacitor: The input current I_i is determined by the input voltage V_i as $I_\mathrm{i} = V_\mathrm{i}/R$; this same current flows in the capacitor, resulting in an output voltage of

$$V_0 = -(1/C)\int I_\mathrm{i}\,dt, \qquad \text{i.e.,} \quad V_0 = -(1/RC)\int V_\mathrm{i}\,dt$$

or

$$V_0 = -(1/t)\int V_\mathrm{i}\,dt \qquad \text{with}\, t = RC$$

Other mechanisms are possible which turn the applied input voltage into a proportional current which is applied to the virtual earth, and they can be thought of as having a "transresistance" equivalent to R in the equation above. The "time constant" of the integrator, t, which can also be thought of as determining its gain, directly determines the frequency behavior of any filter circuit in which the integrator is embedded, and this time constant can be programmed in various ways other than using a resistor; the following subsections describe some approaches.

Continuously Variable Transconductances. These consist of circuits integrated on a single substrate are based on the highly predictable and repeatable physics of the bipolar junction transistor [and also the metal-oxide semiconductor (*MOS*) transistor operating in its weak inversion region] and produce a controllable attenuation, gain, or conductance.

The bipolar junction transistor has a highly predictable relationship between its output current and its input voltage. By relying additionally on the very close matching of devices available on a single substrate, it is possible to produce a range of circuits whose performance can be adjusted accurately enough with externally applied control voltages or currents that they are useful in the construction of programmable filters. The control of such circuits is easily arranged because the production of programmable voltages or currents is quite simple. These circuits tend to be used in the construction of integrators within active filter designs (4).

The OTA. The operational transconductance amplifier (*OTA*) is a voltage input, current output device with a "programming" terminal into which a current is injected which controls the transconductance of the device (5). Filters based on such programmable transconductances are termed g_m–C filters and have become an important technology for the integration of continuous-time filters onto monolithic processes (6,7). OTA integrated circuits are also available, with the classical example being the CA3080 (8). A programmable

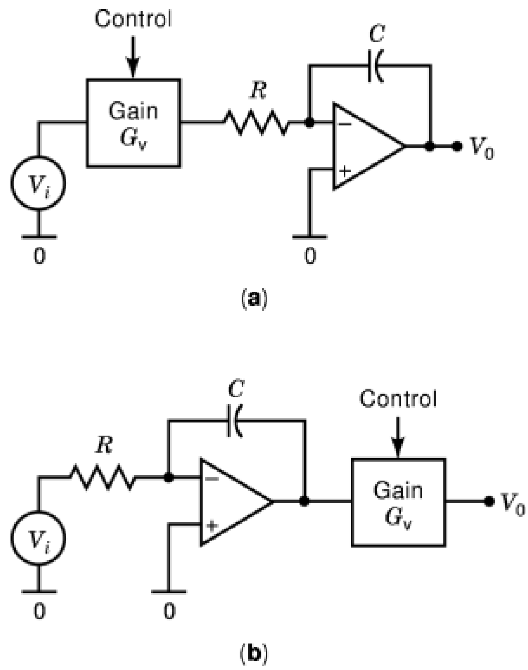

Fig. 4. Two integrators with gain-tuned time constants. Dynamic range considerations determine the choice between them. Any form of gain control element can be used.

integrator is particularly easy to implement with an OTA, requiring simply a grounded capacitor attached to the output.

The VCA or Analog Multiplier. Another device which has found use in programmable filter circuits is the voltage-controlled amplifier (*VCA*). The amplifier can actually be either a voltage amplifier or a current amplifier. A voltage-controlled voltage amplifier with a linear relationship between gain and control voltage is also known as a four-quadrant voltage multiplier; cascading one of these with a fixed time constant integrator is one method of producing an integrator with a voltage-controlled time constant; a careful analysis of dynamic range is required to determine whether the VCA should be located before or after the integrator (Fig. 4). In either case, the effective time constant $t_e = RC/G_v$.

Particularly effective filters can be constructed with the voltage-controlled current amplifier (*VCCA*) (9) which once again finds its place in the integrator, this time shown in Fig. 5. The VCCA is interposed between the resistor R and the "virtual earth" of the operational amplifier around which the capacitor C provides feedback. Impressing an input voltage V_i on the resistor causes a current V_i/R to flow into the VCCA, and the output current I_a (shown in Fig. 5 as directed toward the VCCA) will therefore be $-G_cV_i/R$, where G_c is the current gain of the VCCA. The current flowing in the integrating capacitor is therefore changed by the same amount, and hence the time constant of the integrator is changed:

$$V_0 = -(1/C)\int I_a\,dt, \qquad \text{i.e.,}\quad V_0 = (G_c/RC)\int V_i\,dt$$

so

$$t_e = RC/G_c$$

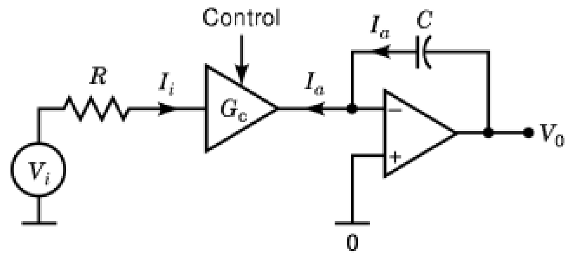

Fig. 5. Noninverting integrator using a voltage-controlled current amplifier.

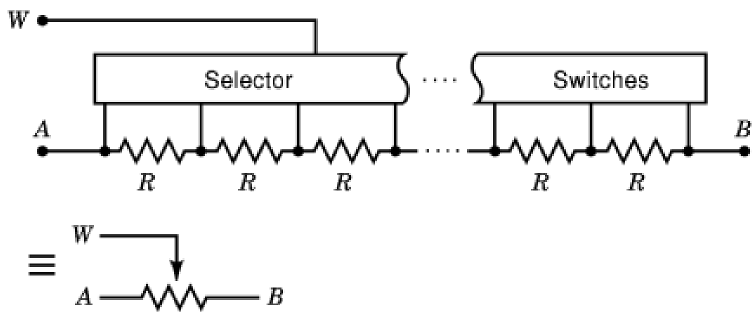

Fig. 6. A digital potentiometer; the resistance of the selector switches does not affect the accuracy of potential division between A and B.

This technique can provide very high linearity performance since there is no signal voltage on either the input or output terminals of the VCCA due to the virtual earth at the input of the operational amplifier—only the current changes with differing input signals.

Switched Voltage or Current Dividers. These are voltage or current dividers and variable transconductances constructed from switched arrays of resistors, whose division ratio can be altered by applying some form of control. This might be thought to be simply an extension of the switched two-terminal component array, but elements which offer attenuation (or gain) are used in quite separate roles in filter circuits from simple two-terminal impedances and are worth considering separately.

Digital Potentiometers. An electronically switched version of the conventional resistive potentiometer has become popular; this version consists of a “string” of resistors connected in series such that each of the points at which adjacent resistors join can be selected for connection to the “wiper” terminal (Fig. 6). This type of circuit is easier to implement in low-cost silicon processes because it is much easier to make a large number of resistors whose absolute value is not too important but whose matching and tracking is quite critical than it is to make even a small number of resistors whose value is accurate and stable. As a result, these “digital pots” find favor in circuits where their rather unpredictable impedance does not affect the circuit response at all, this being done by the very stable division ratio afforded by the chain of well-matched resistors; in other words, these “pots” are being used purely for their division properties and not as variable resistances.

Digital potentiometers are also available with built-in nonvolatile memory, enabling them to store their setting when power is removed. This makes them useful for programming filters where the setting needs to be changed only occasionally (e.g., for calibration purposes).

The R-2R Ladder and the MDAC. Another circuit arrangement of resistors and switching, which can be used to form a programmable divider for either voltage or current, is the “R–2R ladder,” which is a key element of a component called a multiplying digital-to-analog converter (*MDAC*) (10).

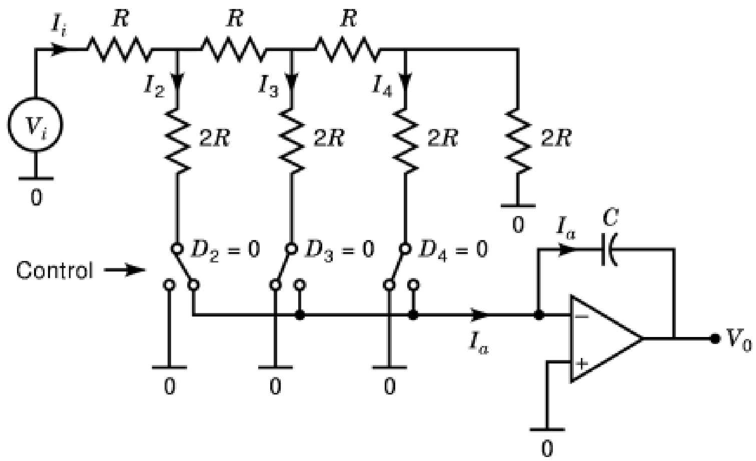

Fig. 7. A 3-bit R–2R network controlling an integrator.

The R–2R ladder has the useful property that, when fed by an input voltage, the voltage on each successive node on the ladder is reduced by a factor of exactly 2 from that on the previous node. To see why this is, we first prove by induction that the input resistance of an R–2R ladder which is terminated at its last node by a resistor of value $2R$ has itself an input resistance of $2R$. Figure 7 shows a single section of the ladder terminated into $2R$; by inspection, the resistance from node 2 to ground is the parallel combination of two $2R$ resistors—in other words, R. The input resistance measured at node 1 is therefore $R + R = 2R$, and the voltage at node two is just $V_i R/(R + R)$ or $V_i/2$.

We can cascade as many of these ladder sections as we require; and note that since the voltage at each successive node is reduced by a factor of 2, then so is the current in each of the $2R$ resistors. In other words, the currents in the grounded arms of the ladder form a binary weighted set. A selection of these binary weighted currents can be routed to the virtual earth of an op-amp instead of to ground, and it can be deployed to program an integrator with a programmable time constant (Fig. 8 shows a circuit with three branches but MDACs are available with up to 16 branches, offering very high resolution):

$$V_0 = -(1/C) \int I_a\, dt$$

with

$$I_a = I_2 D_2 + I_3 D_3 + I_4 D_4$$

where any D_x can equal 0 or 1 depending on whether the branch current is switched to ground or to the op-amp input.

Since $I_i = V_i/2R$ and $I_2 = I_i/2$, $I_3 = I_2/2$, $I_4 = I_3/2$, we have

$$V_0 = -(a/2RC) \int V_i\, dt$$

where here a ranges between 0 and $\frac{7}{8}$.

The fundamental value of the MDAC resistors is sometimes not known to within a certain tolerance error, due to the limitations of integrated circuit (*IC*) processing. The time constant of an integrator made with the MDAC will similarly suffer from this same error because it is proportional to the reference resistance R (11).

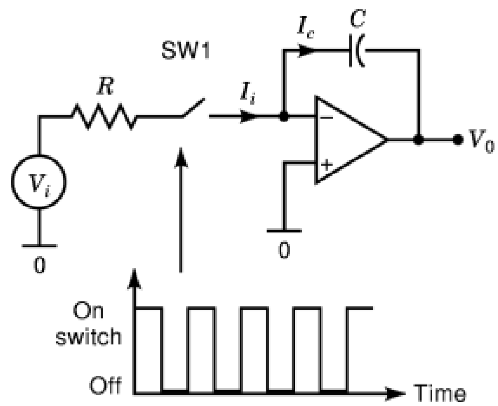

Fig. 8. A mark:space ratio controlled integrator; any suitable electronic switch can be used for SW1.

One solution to this is to use the MDAC instead in a programmable attenuator which can be used to provide an adjustable integrator in the manner of Fig. 4(a) or Fig. 4(b).

Time-Domain Switching. This refers to circuits using some form of signal division in the time domain to manipulate an existing fixed component in order to synthesize a component which appears to have a value which is dependent on the nature of the controlling signal or "clock."

Mark:Space Ratio Control. A technique for varying the apparent value of a resistor is to place it in series with a switch which is periodically opened and closed at a rate which is much higher than the frequency of any interesting signals in the circuit (12). In Fig. 8, if the mark:space ratio of the switching waveform is 1:1, it should be clear that the average value of the current which can flow when a voltage is impressed on the resistor–switch combination is only half that which could flow if the switch were always on. In fact the effective conductance is simply proportional to the fraction m of time that the switch is "on," so $I_i = V_i m/R$. Since this adjustable mark:space ratio can be generated by a very precise digital divider, it is clear that this can be used to turn an accurate resistor into an accurate digitally programmed resistor over quite a wide range:

$$V_0 = -(1/C) \int I_i \, dt, \qquad \text{i.e.,} \quad V_0 = -(m/RC) \int V_i \, dt$$

Switched Capacitor Filters. The other, highly significant technique is that of the "switched capacitor"; this is an important configuration in modern electronics and is covered elsewhere (see Switched capacitor networks). However, a brief review of the basic principle should illustrate how this technique relates to the other methods described here.

Consider Fig. 9, in which the capacitor C_1 is switched between positions 1 and 2 by an input clock, which we shall assume has a 1:1 mark:space ratio. In position 1, the voltage on capacitor C_1 rapidly assumes the value of the input voltage V_i (we assume that the switch resistance is low enough for this to happen). When the switch moves over to position 2, the charge in capacitor C_1 is transferred to C_2; in other words,

$$Q = C_1 V_i$$
$$\Delta V_0 = Q/C_2$$

Therefore

$$\Delta V_0 = V_i C_1 / C_2$$

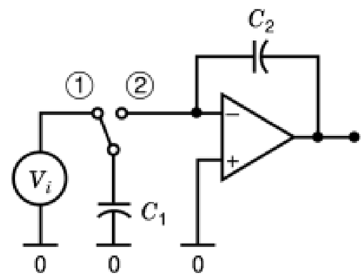

Fig. 9. A switched-capacitor integrator.

Since this transfer happens once every period t where $t = 1/F_s$, we can integrate the successive small changes in V_0 to get

$$V_0 = F_s(C_1/C_2) \int V_i \, dt$$

Once again, this circuit is a controllable integrator, and thus it can be used in any active filter circuit which consists of combinations of integrators—several of which are ideal candidates for programmable filters as we shall now see.

Active Filter Architectures and Programming

This section provides a quick and practically biassed summary. Coverage of particular active filter circuits can be found in Van Valkenburg (13) and particularly in Moschytz and Horn (14). The two most common forms of active filter synthesis are the cascade method and the ladder-derived method; this article will focus on cascade synthesis since it is the most common and practical method for programmable filters.

Filter circuit blocks which have a frequency response which can be described by a pair of poles (and possibly one or two zeroes) are used as basic building blocks to construct a filter circuit of high order; these blocks are called second-order filter sections, or sometimes just “biquads.” Higher-order filter circuits are needed when greater discrimination is needed between different frequencies in the input signal—in common parlance, when a “sharper” filter is required. If the filter circuit blocks can be made programmable by using the techniques of the last section, then several of these blocks can be combined to make a programmable filter with the required response to do the job.

Some Representative Filter Sections and their Applications. A second-order active filter section will consist of at least one amplifier [which is usually an op-amp (voltage in, voltage out)], at least two capacitors, and several resistors. The analysis of the popular sections has been carried out many times, and the results are widely published. Figures 10, 12, and 13 show three second-order all-pole low-pass filter sections in a format representative of published filter textbooks (14), and they show equations which relate their component values and the parameters which the sections are intended to produce. These are expressed here as the pole frequency w_p and the pole quality factor q_p, which is the most common way of describing such a section. Sections which also create a transmission zero are not covered here, but the same principles will apply.

Examples of practical applications are also shown for these filter circuits, but practical considerations such as component choice and the effect of device imperfections are covered in detail in the section entitled “Design for Manufacturability.” Note that there are many more active filter topologies than can be covered by this article, and the reader is encouraged to seek them out and apply these principles to determine how they can be programmed.

Fundamental Restrictions on Programmability. Before examining the individual circuit topologies, it is worth asking the question, "How may components need to be altered in a programmable filter?" These all-pole filters can be completely described by two parameters w_p and q_p, which are functions of the components in the filter. It will usually (but not always) be possible to rearrange the expressions so that two selected component values are a function of all the other component values plus w_p and q_p, which means that in general any filter circuit can be set to the required w_p and q_p by programming just two components. However, the functional relationship between filter parameter and control code may be highly nonlinear, and there may also be forbidden regions where parameter values cannot be set.

Circuits which are useful bases for programmable filters therefore need to conform to more stringent criteria so that the programming approaches described in the previous section can be used effectively. The majority of applications considered require that the pole frequency of the filter section be adjusted while the pole quality factor remains constant. This is needed to ensure that the shape factor of the filter does not change when the frequency is programmed.

This leads to an important rule: If two components are such that pole frequency is a function of only the product of their values and the pole quality factor is a function of only the ratio of their values, then the cutoff frequency of the filter section can be adjusted without change of filter shape by making an equal fractional change to each component. As long as the component ratio stays constant, the pole quality factor does not change. Equally, but of slightly less importance, the quality factor can be altered by changing the ratio of the components, and the pole frequency will not change as long as the product of the component values is invariant.

If, for a given filter section, a pair of components can be found for which the rule above is valid, then the section can be called "canonically programmable." A particular circuit might offer several pairs of components for which it is canonically programmable, in which case the choice depends on the suitability or attractiveness of the programming techniques available for the project.

The VCVS or Sallen & Key Circuit. The first filter section is commonly called the Sallen & Key circuit; requiring only one operational amplifier, set in this case to have unity gain, it has a low component count. Note by examining the equations in Fig. 10 that the pole frequency is a simple function of the four basic passive components, and that changing any one of these will change the frequency. If we look at the expression for the pole quality factor q_p, we can see that this constrains the changes we can make to the components if we also don't want the q_p to change. We can see that q_p is actually a function only of the ratios R_3/R_1 and C_2/C_4, so if we can keep these ratios constant we will keep the q_p constant and thus preserve the shape of the filter; in other words, this section is canonically programmable for (R_1, R_3) and for (C_2, C_4). In this simple Sallen & Key circuit, value-programmable components are ideal control elements, whereas programmable dividers cannot be so easily employed.

In the unity gain lowpass version of the Sallen & Key circuit shown, the resistors can be set to be of equal value whereas the capacitors will be unequal, unless the pole quality factor is equal to 0.5, which is unlikely. Therefore for this lowpass version, it is common to see the tuning carried out by two identical banks of switched resistors—with the capacitors being fixed, or possibly switched in decade steps to give a really large overall range of cutoffs.

The Sallen & Key circuit is also frequently used in a highpass role (Fig. 11), and in this application it is usual to have the capacitors of equal value and use the ratio between the resistors to set the pole quality factor.

The GIC-Derived or Fliege Circuit. The second filter type is usually known as the Generalised Impedance Converter (*GIC*)-derived filter section (Fig. 12 shows the lowpass circuit). Requiring two operational amplifiers and a few more passive components than the Sallen & Key circuit, it has been shown to be much more suited to high values of pole quality factor (15).

Comparing the expressions for w_p and q_p, it can be seen that the circuit is only canonically programmable for (C_1, C_4). This means that the GIC-derived section can be programmed for cutoff frequency with adjustable capacitors. Interestingly, the component R_1 is absent from the expression for the cutoff frequency, and therefore

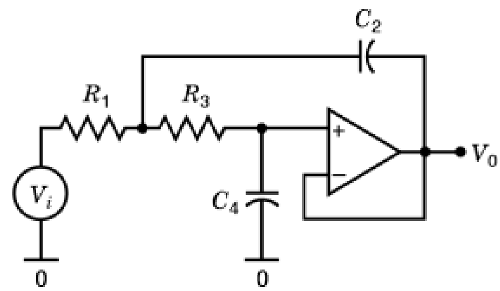

$$T(s) = \frac{w_p^2}{s^2 + (w_p/q_p)\,s + w_p^2} = \frac{V_o}{v_i}$$

$$w_p = \sqrt{\frac{1}{R_1 C_2 R_3 C_4}} \qquad q_p = \sqrt{\frac{R_3 C_2 / R_1 C_4}{1 + R_3/R_1}}$$

Fig. 10. The voltage controlled voltage source (*VCVS*) or Sallen & Key unity-gain low-pass filter. (Adapted from Ref. 14, with permission.)

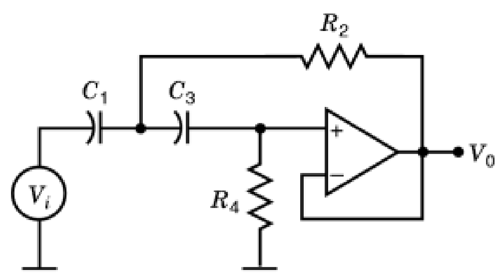

$$T(s) = \frac{s^2}{s^2 + (w_p/q_p)\,s + w_p^2} = \frac{V_0}{v_i}$$

$$w_p = \sqrt{\frac{1}{C_1 R_2 C_3 R_4}} \qquad q_p = \sqrt{\frac{R_4 C_1 / R_2 C_3}{1 + C_1/C_3}}$$

Fig. 11. The Sallen & Key unity gain high-pass filter. (Adapted from Ref. 14, with permission.)

it is possible to adjust the pole quality factor freely by adjusting this resistor, which could be implemented as a single programmable resistor array.

Variations of this circuit cover high-pass, bandpass, and notch responses and share the same characteristic as the lowpass circuit, namely that the pole frequency can only be canonically tuned by adjusting the capacitor values.

State-Variable Circuits. The third filter type, shown in Fig. 13, is variously known as the state-variable or Kerwin–Huelsman–Newcombe (*KHN*) biquad. It employs three operational amplifiers, which might be considered to be wasteful in comparison to other circuits, but the cost and area penalties using modern amplifier technologies are small. The only significant disadvantage is the increased power consumption entailed.

Inspection once more of the expressions for w_p and q_p shows that both the resistor pair (R_5, R_6) and the capacitor pair (C_6, C_8) meet the criterion for canonical programmability. The ratio of resistors R_3 and R_4

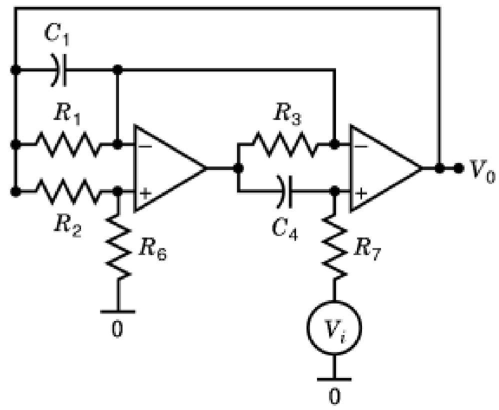

$$T(s) = \frac{Kw_p^2}{s^2 + (w_p/q_p)\,s + w_p^2} = \frac{V_0}{v_i}$$

$$K = 1 + \frac{R_2}{R_6}$$

$$w_p = \sqrt{\frac{R_6}{R_2R_3R_7C_1C_4}}$$

$$q_p = R_1\sqrt{\frac{R_6C_1}{R_2R_3R_7C_4}}$$

Fig. 12. The generalized impedance converter (*GIC*)-derived or Fliege low-pass filter. (Adapted from Ref. 14, with permission.)

appears in both expressions, meaning that this ratio should best be left alone. The ratio of resistors R_1 and R_2 appears only in the expression for q_p, providing an independent method for adjusting the pole quality factor, which will be found useful.

Examining the topology more closely, it can be seen that the component pairs (R_5, C_6) and (R_7, C_8) are each employed with an amplifier to form an integrator; in fact, only the products R_5C_6 and R_7C_8 appear in the expressions for w_p and q_p. In other words, the structure is canonically programmable for integrator time constants (t_1, t_2). This is extremely useful because it means that the integrators can be replaced with any of the programmable integrator structures discussed in the previous section. This versatility has resulted in the KHN biquad being the filter topology of choice for the majority of programmable filter applications requiring high performance and closely predictable responses.

Most of the programmable integrator techniques discussed earlier can be employed in a KHN biquad circuit. Figure 14 shows such a biquad programmed by binary-weighted switched resistors and decade-selected capacitors; all the switching is done by an integrated switch array. Note that the switches are referred to the inverting inputs of the integrator amplifiers, which means that there is no signal swing on a conducting switch (see section entitled "Component Nonlinearities and Imperfections"). It is important to ensure that enough control voltage is available to turn the switches on and off fully over the entire signal voltage swing; commercially available integrated complementary metal-oxide semiconductor (*CMOS*) switch packages will achieve this, but more care must be taken with individual *JFET* or DMOS switches.

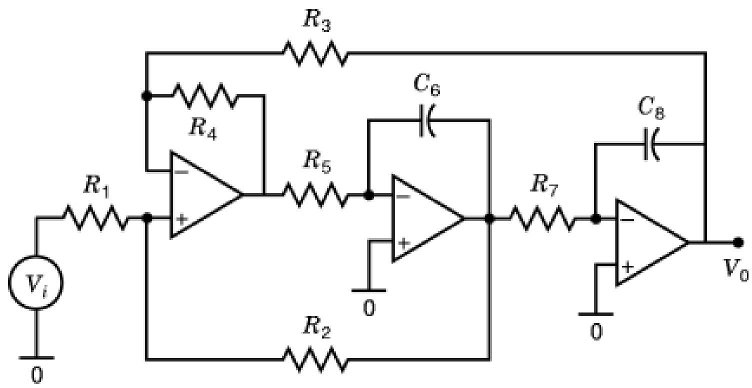

$$T_{LP}(s) = \frac{K_{LP} w_p^2}{s^2 + (w_p/q_p)\, s + w_p^2} = \frac{V_0}{v_i}$$

$$K_{LP} = \frac{R_3 R_4 + 1}{R_1/R_2 + 1}$$

$$w_p = \sqrt{\frac{R_4}{R_3 R_5 R_7 C_5 C_8}}$$

$$q_p = \frac{1 + R_2/R_1}{1 + R_4/R_3} \sqrt{\frac{R_4 R_7 C_8}{R_3 R_5 C_6}}$$

Fig. 13. The KHN state-variable low-pass filter, shown in its noninverting configuration. (Adapted from Ref. 14, with permission.)

Figure 15 shows the classical circuit for using an OTA in a KHN biquad circuit (8). This circuit is limited in its dynamic range not by fundamental constraints but because the integrated circuit OTAs currently available have a very limited input voltage range, only a few tens of millivolts for linear operation, which is the reason that input attenuation must be used at the front of each OTA.

An alternative topology (Fig. 16) using OTAs relies on their very high output impedance to allow a grounded capacitor connected at the output to directly integrate the output current. The capacitor is then buffered to prevent the next stage from loading it. Integrated circuits combining OTAs and buffers are readily available. In this circuit the integrators are noninverting, which allows a passive summation network to be used to combine the state feedback, eliminating one op-amp.

A common audio frequency application of the programmable state-variable circuit is in the design of equalizers for audio systems. Reference 9 describes a representative circuit employing VCCA tuning for the integrators in a KHN biquad; the current actually injected into the virtual earth of the amplifier is equal to the current flowing through the VCCA's input resistor multiplied by a scaling factor which is the logarithm of the input control voltage divided by a reference voltage. This results in a linear decibels per millivolt function which is appropriate for audio applications. Note also that the VCCA used here produces an output current which is in the same direction with respect to the VCCA as the input current is. From the point of view of the filter circuit, this turns the integrator in which the VCCA is embedded into a noninverting one, and so the polarity of the feedback from the output of the first integrator must be reversed to allow for this. The reference

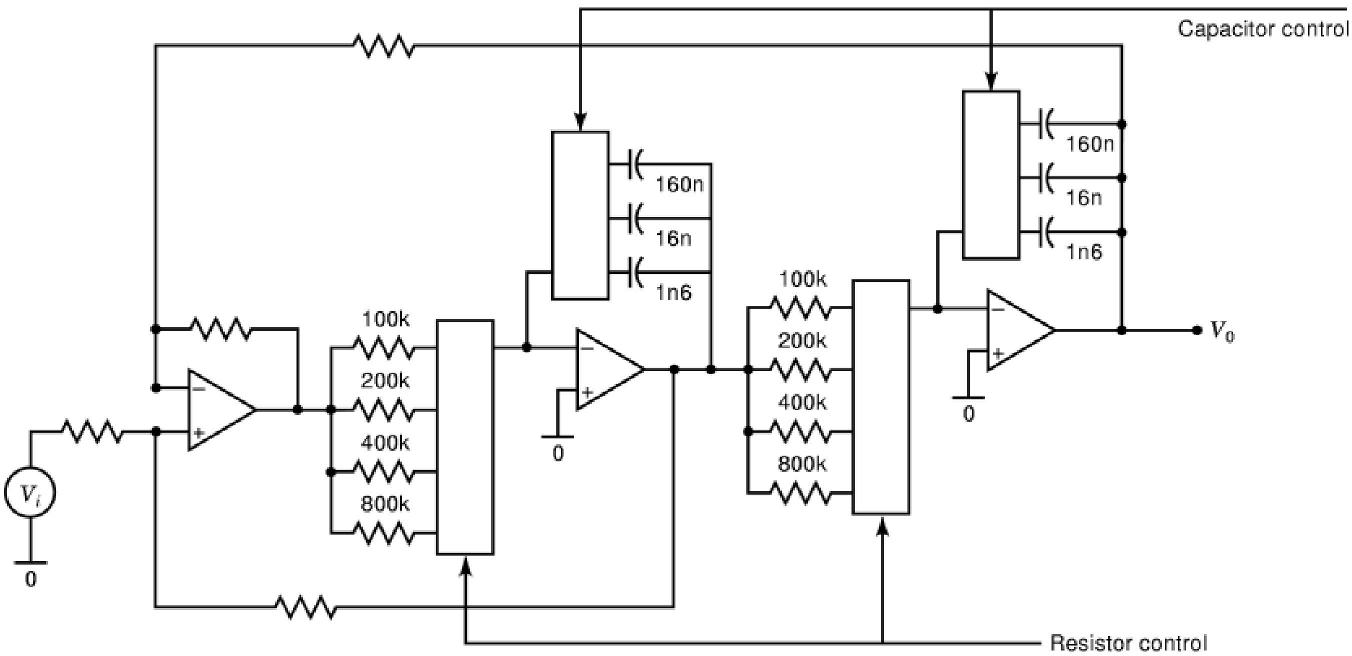

Fig. 14. The KHN low-pass filter programmed with arrays of resistors and capacitors.

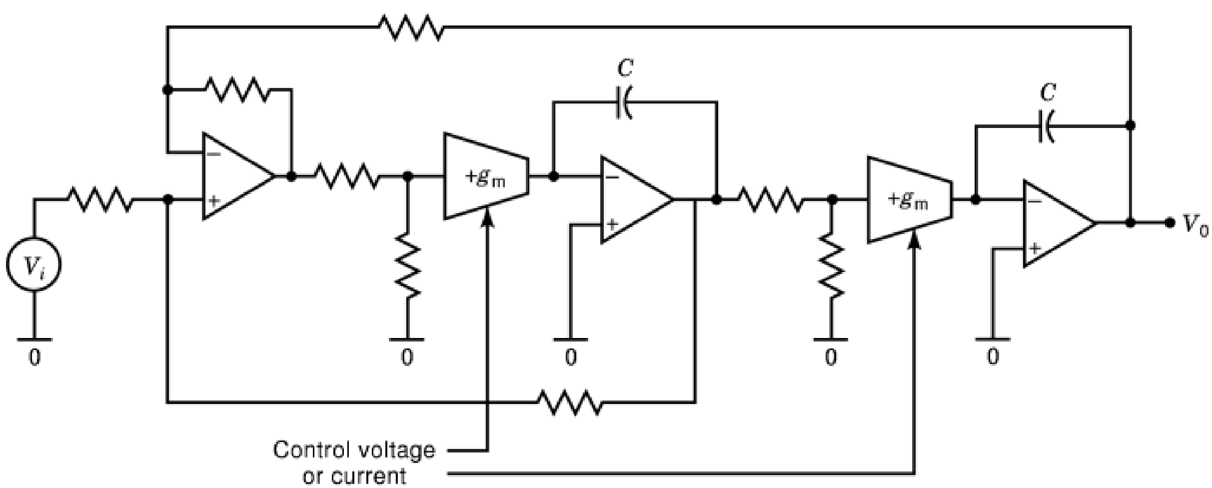

Fig. 15. A KHN low-pass filter with op-amp integrators tuned by operational transconductance amplifiers (OTAs).

also illustrates the ease in which a VCCA-based circuit can be configured to provide independent control over w_p and q_p in this topology.

MDAC-Programmed State-Variable Filters. The MDAC has found wide application in programmable KHN biquads; many manufacturers produce highly accurate MDAC devices and provide application notes on their use in filter design. Figure 17 shows a circuit in which not only the integrator time-constants but also the feedback paths which were represented by (R_3, R_4) and (R_1, R_2) (see Fig. 13) are replaced with MDACs (11),

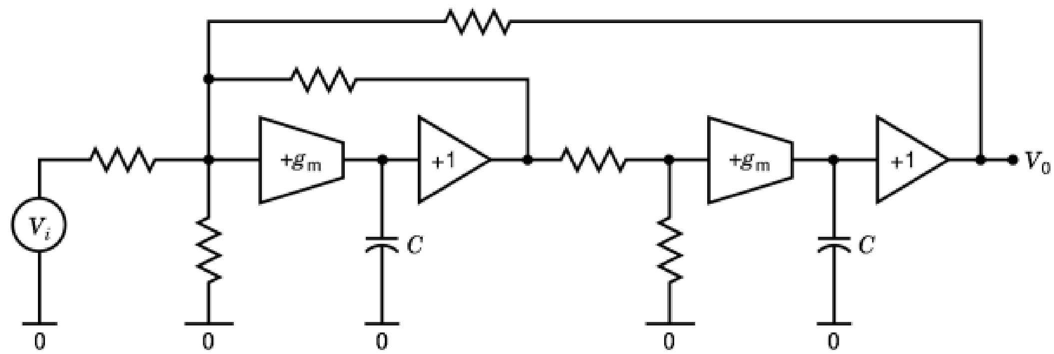

Fig. 16. An alternative OTA-tuned low-pass filter with two integrators, not requiring a summing amplifier.

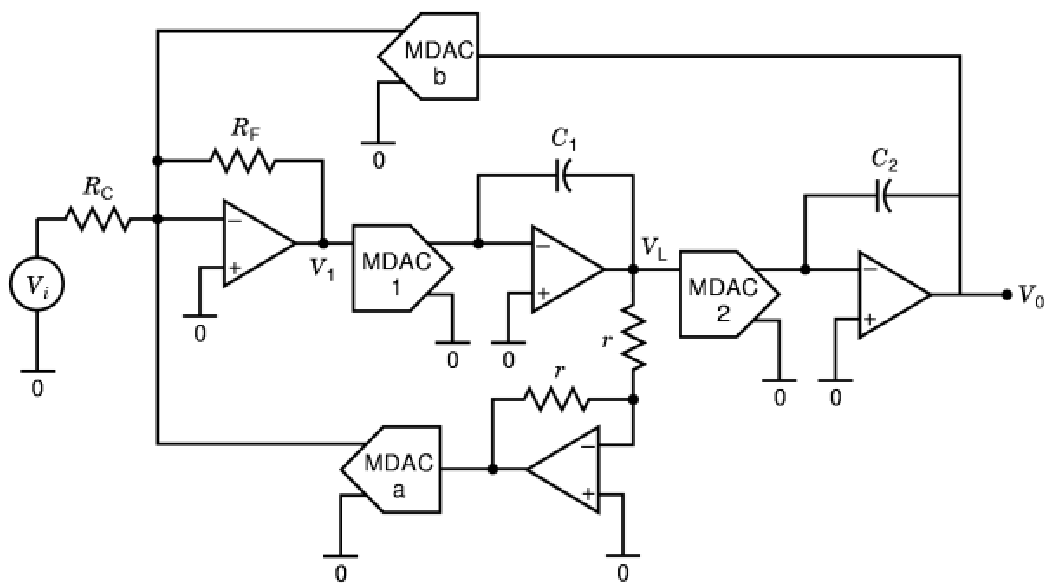

Fig. 17. An MDAC-programmed KHN state-variable low-pass filter, with additional MDACs to give finer control of pole frequency and quality factor.

with an extra inverter employed to ensure that all feedback can be returned to the virtual earth of the input op-amp. The extra MDACs a and b provide higher resolution for setting the w_p and q_p values of the section; the MDACs controlling the integrators (in multiple biquads which may be cascaded to make a higher order filter) can be controlled with a common code. The tolerances of the programming MDACs' fundamental resistance can be compensated for with the extra MDACs.

The MDAC-programmed integrator has several subtleties which can affect the performance of filters into which it is embedded. Firstly the output capacitance is a nonlinear function of the code applied, which can cause the behavior of filters using this technique to depart from ideal at high frequencies (16).

A further difficulty is not immediately obvious but restricts the performance of MDAC-based filters primarily at the lower extremes of their programming range and also affects any other tuning method in which some form of resistive voltage or current division is used to alter the integrator time constant. The operational amplifiers used to construct the integrators contribute to the overall noise and direct-current (*dc*)

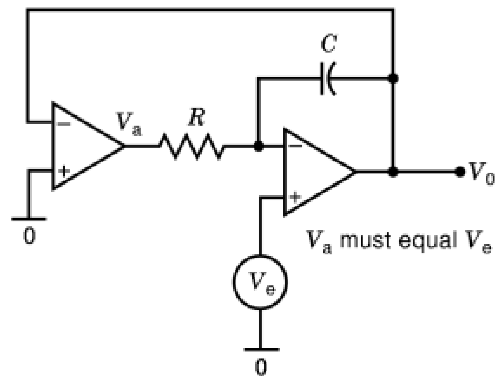

Fig. 18. A test circuit to explore the effect of the dc offset of a standard integrator in a feedback configuration.

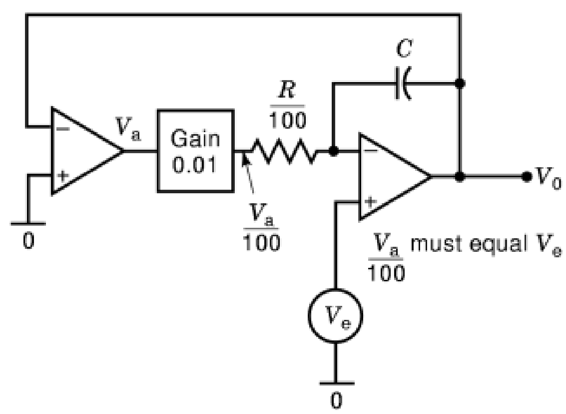

Fig. 19. A test circuit to explore the effect of the dc offset of an attenuator-tuned integrator in a feedback configuration.

offset performance of the filters; and the greater the voltage division set by the programming circuits for the integrators, the greater is the consequent magnification of these effects, due to the overall dc feedback which exists around the KHN circuit (17).

Consider the two integrators of identical time constant as shown embedded in the circuits of Figs. 18 and 19. In Fig. 18 the time constant is set to be *RC* by means of components of those values; in Fig. 19 the same time constant is achieved by feeding a resistor of value $R/100$ through an attenuator with a gain of 0.01; this produces the same current into the integrator and hence the same time constant. These integrators are each embedded in the feedback loop of an ideal test op-amp with grounded input to force the integrator output V_0 to zero (note that feedback is returned to the noninverting input of this test amplifier since the integrators invert; also note that the op-amps in the integrators are assumed ideal for the moment). The noninverting input of the integrator op-amps would normally be grounded, but let us inject a small voltage V_e into each inverting input and work out what the change in the output voltage of the test op-amp will be. In the case of Fig. 18 we can see that to balance the voltage V_e the output of the test opamp V_a will equal V_e; at dc the capacitor has no effect and can be ignored. This is because there is no voltage dropped across the resistor *R*.

In Fig. 19 we can see by the same argument that the output of the attenuator has to be equal to V_e to balance the dc conditions, but this means that the output of the test amplifier is now $100V_e$ to achieve this.

These test circuits are not real filter circuits, but the split feedback loops in a state-variable filter mean that the effective attenuation caused by either an MDAC or for instance a "digital pot" causes a buildup of dc offset, and also low frequency noise, related to the input behavior of the integrator op-amps; it worsens as the effective attenuation factor increases—that is, toward low cutoff frequencies. This effect is not seen when programmable resistances are used, and it allows the latter type of filter to offer superior stability and dynamic range compared to MDAC-tuned filters using the same filter op-amps.

Design for Manufacturability

Programmable filters of the types referred to here tend to be manufactured from separate general-purpose electronic components, although some manufacturers have produced integrated circuits combining several of the required functions (18). The integration of entire circuits of this type onto a monolithic substrate presents many problems, which is why alternative techniques such as switched capacitor filtering have become popular. This section highlights issues relating to the use of real components in programmable filters.

Device Tolerances. Clearly, since the frequency response of a filter circuit is a function of the values of the components in it, changes in these values will change the filter response. For some special applications, it is possible to have resistors and capacitors manufactured to the exact values determined by theory; but this is impractical for most filter designs, whether fixed shape or programmable. Normally, the designer must recognize the availability of "preferred value ranges" in the design process. The finite tolerances and the availability of only discrete values of components impacts the design of programmable filters employing directly switched components; the small deviations of the values of these components from ideal will cause the frequency response of the realized filter, after it has been normalized to the intended cutoff frequency, to vary. Controlling this variation usually leads to the requirement for closely matched and toleranced components. In this respect, control techniques using MDACs, which are trimmed very accurately during the production phase, can be very helpful.

Op-Amp Bandwidth and Predistortion. Predistortion is the term applied to the alteration of the values used in a filter circuit so that the finite bandwidth effects of the amplifiers used are exactly compensated for. If a programmable filter needs to cover a wide frequency range, it is quite possible that it will have to work in regions where predistortion would not otherwise be an issue, as well as in regions where the "undistorted" values would produce an incorrect response due to amplifier limitations. The presence of this problem presupposes that it has not been possible within the design to find an amplifier whose performance is in all respects adequate, and that some form of compensation will be required.

There are two approaches to a resolution of this problem. Firstly, it may be possible to set the adjustable components or circuit elements (e.g., integrators) to values which are different from what would be expected if the amplifiers were perfect, but in such a way that both the pole frequency and the pole quality factor are achieved. This will require higher resolution from the controlled circuit elements than would be needed simply to achieve the individual frequency steps required. For instance, a pair of simple 4-bit binary-switched resistor arrays would normally be sufficient to provide a 15:1 frequency adjustment for say a Sallen & Key circuit or a KHN Biquad at lower frequencies, but at high frequencies the 15 different values needed by each resistor (and they will probably not be equal) might require at least 8-bit resolution, immediately doubling the complexity of the control elements (19).

In fact, as mentioned in the section entitled "Parasitic Capacitance," this increase in the complexity of the control elements may itself inject another form of variability into the parameter setting process due to the variation of parasitic capacitances which cannot be disregarded at higher frequencies, and which can combine to make the pole frequency and pole quality factor an almost unpredictable function of the digital control word, making automated functional calibration on a computer-controlled test fixture the only realistic way of manufacturing these circuits (16).

However, the use of voltage-controlled circuit elements blends in well with this approach, since it is usually quite easy to provide large numbers of static or slowly varying dc control signals from the central control processor, often using a Pulse Width Modulation (*PWM*) technique which requires little in the way of expensive, high-accuracy DAC components.

The other way of resolving the problem is to find a compensation technique for the circuit topology which is orthogonal to the process for generating the theoretical component values. This involves additional (usually passive) circuit elements which are added to the circuit so that, to at least a good approximation, the circuit produces the intended response parameters when a particular set of component values, intended for use with ideal amplifiers, are selected by the control circuitry.

By far the most important example of this technique is the use of phase-lead compensation in the two popular biquads based on a pair of integrators: (a) the KHN biquad already discussed in detail, and (b) its close cousin the Tow–Thomas biquad (20). In this technique a fixed capacitor is placed across one of the resistors in the filter which is not adjusted when the filter is used in a programmable context (nearly always R_3 in Fig. 13); the *RC* time constant is set to provide phase lead which cancels out the phase lag introduced by the three amplifiers. This compensation is effective whatever the components used in the two integrators, and it is a sound practical way to compensate programmable KHN filters. To a first approximation, the value of the time constant for such a compensation capacitor used in a loop with three amplifiers, two of which are realizing a standard integrator, is

$$t_c = (2 + \text{ the noise gain of the summing amplifier})/w_a$$

where w_a, the unity gain frequency of the amplifiers, is measured in rads/second. If we know the value of the resistor, the capacitor value follows directly.

Parasitic Capacitance. Layout parasitics are generally constant and unrelated to the activities of the control circuits in a programmable filter circuit. Such parasitics often affect the performance of an active filter and should be taken into account in the design. Layout parasitics can be expanded to include the effects of input capacitance of the amplifiers used; this is a significant source of error in many circuits and must also be considered. Naturally, filter circuits containing only grounded capacitors are the easiest to make allowances for (since the parasitic input capacitance of the amplifier is effectively connected to ground because ground and the power supplies are, or should be, common at high frequencies).

However, a separate class of parasitics plague programmable filters and can be the limiting factor in their successful operation at higher frequencies. All programming techniques which deploy an array of switches to link components into the circuit—and this includes the MDAC structures discussed earlier—are affected by the variation of the parasitic capacitance of the switch elements between the "off" and "on" states. In addition, since various circuit branches are switched in and out by this process, static parasitics associated with these branches are switched too, making the total amount of parasitic capacitance a complex function of the digital code used to control the filter parameter.

The reason why this causes an unwelcome effect rather than just a smoothly varying degradation which can be coped with in a similar way to the problem caused by op-amp bandwidth limiting relates to the way in which a wide range of component value is implemented efficiently using some form of binary coding. As seen earlier, this applies to both binary-weighted component networks and to MDACs.

Consider the integrator circuit of Fig. 20, in which a 4-bit binary controlled resistor network is being used to control the time constant. The relevant parasitic capacitances are also included; C_1 through C_4 are the stray capacitances across the resistor elements (note that these will in general not be related to the resistance value in any simple way) while the switch terminal capacitances are shown as C_i and C_0 (again, note that these capacitances depend on the state of the switch). As shown previously, this circuit is capable of adjusting the cutoff frequency of a filter in which this integrator, and others like it, are embedded, between F and $15F$, where

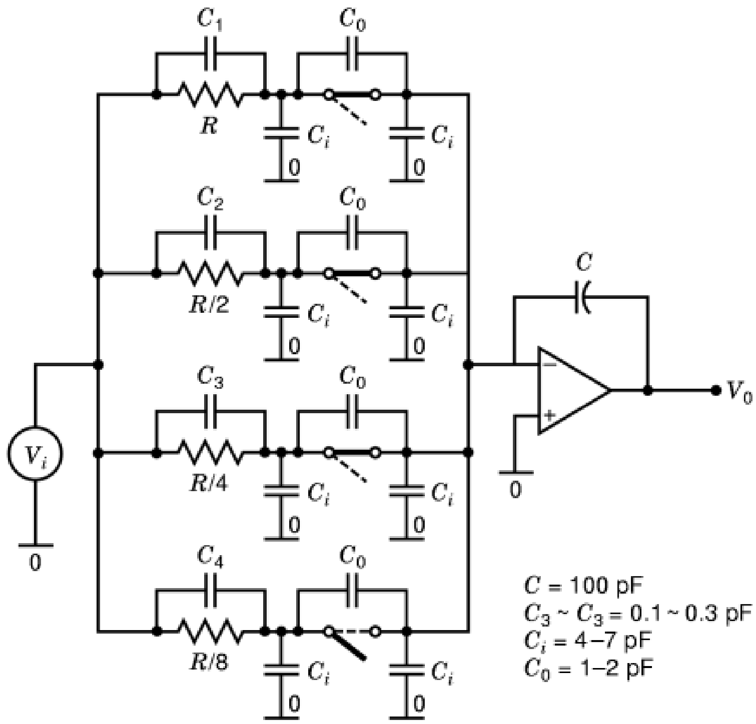

Fig. 20. A resistively programmed inverting integrator showing the parasitic capacitances associated with typical 4-bit control.

F will be set by all the other components in the filter but will be proportional to $1/RC$. With switches in the solid position, the cutoff is being set to $7F$; the most significant bit is off and all the rest are on. With switches in the dotted position, the cutoff is being set to $8F$ and the most significant bit is now on, with the rest off. In one case, we have three "on" switches, together with their own parasitics and those of the connected sub-branches (the resistors themselves), while in the other case, only one switch is on. In particular the effect of the stray capacitance on the resistors causes a variable degradation of the performance of the integrator.

Since there is a very nonlinear relationship between a number and the number of "on" binary bits needed for its representation, then if the parasitic elements in a circuit are significant, their effect will vary in a very complex way as the filter cutoff frequency is programmed using a switched resistor technique. Solutions to this problem generally involve ensuring that the ratio between the fixed capacitor used in the integrator and the parasitic capacitances introduced by the switched components is as high as possible.

Component Nonlinearities and Imperfections. The effect of amplifier nonlinearity on filter performance is common to both programmable and fixed response filters; however, the particular techniques used to achieve programmability have their own influence on nonlinearity and also to a certain extent on the noise developed by the circuit.

Circuits using electronic switches will be sensitive to the nonlinear behavior of the on resistance of the switch used, as a function of signal voltage swing. Improving process technologies bring steady reductions in this source of distortion, but clearly the best solution is to eliminate the signal voltage swing by some means. The KHN biquad employing switched resistors (Fig. 14) achieves this by operating the "on" switches at a virtual earth point, so that whatever the signal swings at the outputs of the amplifiers in the filter, there is no modulation of the on resistance.

When analog switches are used to switch banks of capacitors into a filter circuit, the on resistance of the switch will cause an error in the filter response which is difficult to compensate for. This resistance should be kept as low as possible; if space and power consumption are not critical, small relays can be used instead of silicon-based switches.

Even when used in series with an intentional resistor, the on resistance behavior of a silicon switch can cause problems due to the high positive temperature coefficient of resistance which it displays. The effect of this resistance temperature coefficient should be taken into account; it makes little sense to spend money on ultrastable resistors only to throw it away with this source of error (17).

Testing. Filters of this complexity present a greater alignment and testing challenge than those with a fixed shape and cutoff frequency, because the programmability must be confirmed for the whole range over which it is offered. It may be that a single specification for the realized performance of the filter is applied whatever the programmed parameter, or the deterioration of performance occurring toward the extremes of the range of programmability may be recognized by a loosening of tolerances at these extremes.

Automated testing in which both the programmable filter under test and the test equipment employed are sent through a sequence of settings is essential for testing to be both efficient and effective. The nature of the tests employed can be thought of as providing a “test vector” which is quite analogous to that employed in the testing of digital circuits. The test vector, or set of tests to be performed, must be constructed to ensure that (1) the end user of the circuit can see clearly that specification-related tests have been performed and have been passed and (2) any potential areas known to the designer where failure may occur have been carefully explored.

BIBLIOGRAPHY

1. U. Kumar A decade range linearly tuneable versatile filter, *IEEE Trans. Instrum. Meas.*, **IM-30**: 165–167, 1981.
2. E. Damm Spannungsgesteuerte Analogfilter, *Elektronik*, **10**: 90–92, 1976.
3. A. M. Durham W. Redman-White J. B. Hughes Digitally-tuneable continuous time filters with high signal linearity, *IEE Colloquium Digital Analogue Filter Filtering Syst.*, 1991, pp. 9/1–9/4.
4. R. Schaumann Design of continuous-time fully integrated filters: A review, *IEE Proc.*, **136** (4, Pt. G): 184–190, 1989.
5. S. Franco Use transconductance amplifiers, *Electron. Des.*, **24** (19): 98–101, 1976.
6. Y. P. Tsividis J. O. Voorman (eds.) *Integrated Continuous-Time Filters: Principles, Designs & Applications*, New York: IEEE Press, 1993.
7. K. Fukahori A bipolar voltage-controlled tuneable filter, *IEEE J. Solid-State Circuits*, **SC-16**: 729–737, 1981.
8. H. A. Wittlinger *Applications of the CA3080 high performance operational transconductance amplifier*, application note ICAN6668, NJ: RCA, 1973.
9. G. K. Hebert F. Floru Digitally-controllable audio filters and equalizers, *13th Int. Conf. AES*, Dallas, TX, 1994.
10. D. J. Sheingold *Application Guide to CMOS Multiplying D/A Converters*, Norwood, MA: Analog Devices, 1978, pp. 12–14.
11. J. M. Zurada Digitally tuneable active filter biquads, *Proc. Southeastcon ’81*, 1981, Huntsville, AL, pp. 364–368.
12. R. W. Harris H. T. Lee Digitally controlled, conductance tuneable active filters, *IEEE J. Solid-State Circuits*, **SC-10**: 182–185, 1975.
13. M. E. Van Valkenburg *Analog Filter Design*, Philadephia: Holt-Saunders, 1982.
14. G. S. Moschytz P. Horn *Active Filter Design Handbook*, New York: Wiley, 1981.
15. N. Fliege A new class of second-order RC-active filters with two operational amplifiers, *Nachrichtentech. Z.*, **26** (6): 279–282,1973.
16. K. Goodman J. M. Zurada Frequency limitations of programmable active biquadratic filters, *Proc. Southeastcon ’81*, 1981, Huntsville, AL, pp. 369–373.
17. K. Castor-Perry Using resistor-set filters, application note 7, Kemo, Ltd.
18. UAF42 Universal active filter data sheet, Tucson, AZ: Burr-Brown.

19. K. Goodman J. M. Zurada A predistortion technique for RC-active biquadratic filters, *Int. J. Circuit Theory Appl.*, **11**: 289–302, 1983.
20. L. C. Thomas The biquad: part I—some practical design considerations, *IEEE Trans. Circuit Theory*, **CT-18**: 350–357, 1971.

READING LIST

D. Lancaster *Active Filter Cookbook*, Indianapolis, IN: Sams.

P. Horowitz W. Hill *The Art of Electronics*, 2nd ed., Cambridge, UK: Cambridge Univ. Press, 1989.

R. Schaumann M. A. Soderstrand K. R. Laker *Modern Active Filter Design*, New York: IEEE Press, 1981.

W. Heinlein H. Holmes *Active Filters for Integrated Circuits*, New York: Springer-Verlag, 1974.

KENDALL CASTOR-PERRY
Burr-Brown Corp.