Take tunable lowpass filters to new heights

TOM NAPIER, CONSULTANT

Lowpass filters separate a desired signal from a higher frequency, unwanted signal; remove residual carrier from a demodulated signal; re-move high-frequency noise; and even filter a signal so that subsequent sampling generates no in-band aliases. Often, you can use a fixed-frequency filter, but sometimes you need to tune the filter to suit the input signal, either because the system must operate with a range of inputs or because the characteristics of the signal itself may change.

With current output multipliers and fast amplifiers, you can design active lowpass filters whose cutoff frequencies you can electrically tune from 10 Hz to at least 10 MHz.

Using a tunable active filter may be the best design strategy. Traditionally, active filters have been applied below 100 kHz. However, their range of operation can extend to at least 10 MHz.

Filters start with a little theory

The classical Sallen-Key active filter uses two resistors, two capacitors, and a buffer amplifier to implement a two-pole filter (references ` and ´). You can implement even-order, all-pole filters as a series of Sallen-Key stages, adding an extra RC stage if you need an odd number of poles. In practice, the high Q of some pole pairs and the need for accurate components make active filters with more than eight poles tricky to implement. However, two- and four-pole filters suffice for many applications.

The Sallen-Key filter has a cutoff frequency that is a function of its two—usually equal—resistors, and you can manually tune the filter with a dual potentiometer. Alternatively, you can switch capacitors to increase the tuning range. Electrical tuning is more difficult because it requires matched, isolated, controlled resistors. At one time, cadmium-sulfide photoresistors with built-in incandescent bulbs or LEDs were used for electrical tuning, but these photoresistors are now hard to find. Another practical approach to single-range filters uses fixed resistors and varactor diodes, and an honorable mention should also go to designs that use multiplying DACs to implement digital tuning (reference ..).

The other tunable active-filter configuration you can use is the state-variable filter or its variant, the bi-quad filter. The state-variable filter combines two integrators and a summing amplifier into a two-pole filter. The output of the summing amplifier is a highpass-filtered version of the input, and the first integrator supplies a bandpass output. These outputs may be useful in their own right, for example, for synthesizing allpass and band-reject filters. You tune the state-variable configuration by varying two resistors.

By putting an analog multiplier and a fixed resistor in front of each integrator, you make the state-variable filter electrically tunable. Such filters, which appeared in the mid-1970s, require three amplifiers per two-pole section (reference _). This filter topology also has the disadvantage that one of its feedback loops passes through two amplifiers, and the other loop passes through three amplifiers. This configuration makes the filter sensitive to excess phase shift that results from stray capacitance and limited amplifier bandwidth. Thus, early state-variable filters didn't tune much above 100 kHz.

With the advent of analog multiplier ICs having differential inputs and current outputs, designers were able to make tuned filters without a summing amplifier. The most practical configuration is based on the Sallen-Key filter (ϕigure `). This configuration doesn't offer the state-variable filter's bandpass output, but it is more stable than the more conventional designs at high frequencies. With care, you can push the filter's design to operate at as much as 50 MHz.

You can linearly tune this filter by applying a control voltage,

FIGURE 1

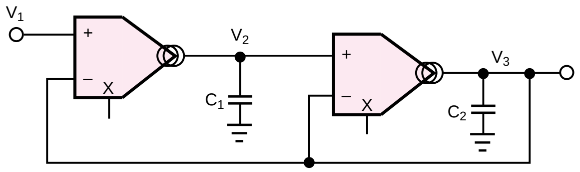

$V_3/V_1=1/(1+sRC_1+s^2R^2C_1C_2)$.

$V_2/V_1=(1+sRC_2)/(1+sRC_1+s^2R^2C_1C_2)$.

This Sallen-Key filter is implemented with current-output multipliers, where R is the multipliers' voltage-controlled transimpedance.

X, to each multiplier. Varying this voltage changes the effective value of the multiplier's transfer resistance, R. With this technique, you can easily tune the filter stage over a 10-to-1 frequency range, and 50-to-1 is possible. You can achieve an even wider range by switching the integrating capacitors in decades using either relays or solid-state switches.

If you use a multiplier IC that has a high output compliance, such as the HA-2547 from Harris (Melbourne, FL), place the integrating capacitors between the multiplier outputs and ground exactly as shown in ϕigure `*. This placement simplifies range switching. Because the Harris ICs have a high input impedance, you can connect feedback from the outputs to the inputs without loading the outputs. Only the output driving the rest of the system requires a buffer amplifier, and the result is a two-chip tunable filter (*ʀeference *).*

Limits of the simple filter

One disadvantage of the circuit in ϕigure ` *is that the output capacitance of the multiplier, which is 6 to 10 pF, is part of the integrating capacitance. Whenever you need a tuning capacitance smaller than about 100 pF, the multiplier capacitance becomes a large and uncertain part of the tuning equation. This multiplier capacitance also limits the highest cutoff frequency because the other factor of the RC term is the transfer impedance of the multiplier. In the HA-2547, for example, this impedance is variable down to 2500γ. Thus, the minimum RC time constant is about 25 nsec.*

*A partial solution to the capacitance problem is to add a conventional integrating amplifier to each multiplier output and reverse the sense of the feedback connections to compensate for the amplifier inversion (*ϕigure ´*). With this approach, you can neglect the multiplier output capacitance to a first approximation, and, with a fast voltage-feedback amplifier, the integrating capacitor can be as low as 5 pF. The EL2073C amplifier from Elantec (Milpitas, CA), for example, works well in this application. Don't be tempted to use a current-feedback amplifier here, because these devices are unstable with capacitative feedback.*

Another disadvantage of the configurations in ϕigures ` *and* ´ *is that unless you apply different control voltages to the two multipliers, the Q of the filter is proportional to the square root of the ratio of the two integrating capacitors. When you require either a high or a low Q, the ratio of these capacitors becomes large, which again limits the highest cutoff frequency that the filter can reach.*

FIGURE 2

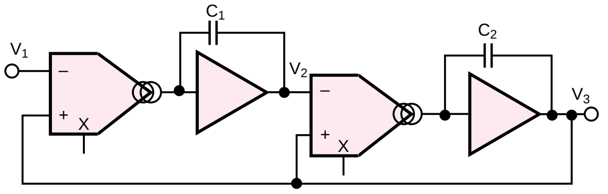

Adding integrating amplifiers reduces the effect of the multiplier capacitance, allowing a higher cutoff frequency. Note the reversed multiplier input connections. The transfer function is the same as in Figure 1.

FIGURE 3

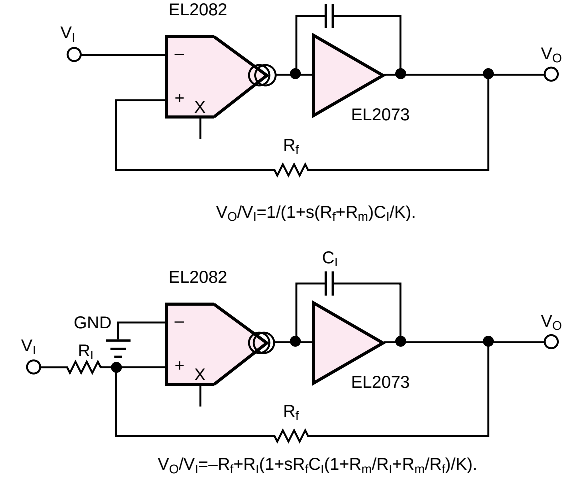

Single-pole filters using an Elantec EL2082 multiplier offer an alternative topology. K is the multiplier's current gain; R_m is its input resistance, about 95 .

These filter configurations have an inherent dc gain of unity, guaranteed by the amplifiers' dc open-loop gain and the dc feedback from the output to the input. However, high-frequency amplifiers tend to have an open-loop gain of less than 1000 and a relatively low input impedance. If you try to lower the filter bandwidth by applying too low a voltage to the multiplier, the loop gain falls and the dc gain is less than unity. If high-frequency performance is unnecessary, it pays to use a slower, higher gain amplifier.

The HA-2547 multiplier has an inherent bandwidth around 40 MHz and works well at filter-cutoff frequencies below 1 MHz. It has enjoyed use in successful commercial products at cutoff frequencies as high as 5 MHz. However, to reach higher frequencies, either smaller capacitors or a lower transfer impedance is necessary.

Pushing beyond the limits

When you try to implement a filter with a cutoff frequency higher than 5 MHz, the Elantec EL2082 multiplier IC is a good choice. This device differs considerably from the HA-2547 multiplier. It has eight rather than 16 pins, which makes the design more compact. More important, you can tune its transfer impedance as low as 50γ.

The signal input for the Elantec multiplier is a current sink with an impedance of 95γ, rather than the HA-2547's differential voltage input. However, you can still use this input differentially, because the chip has a high-impedance input reference pin that drives the other end of the 95γ input impedance.

If a pure current source drives the multiplier's input reference, *then the voltage on the reference pin has no effect. However, if a voltage source drives the input via an external resistor, the input current is a function of the difference between the source voltage and the reference voltage. This situation allows you to construct many interesting filter configurations.* ϕigure .. *shows two ways of implementing a single-pole filter, and* ϕigure _ *shows three practical two-pole-filter configurations based on combinations of the sin- gle-pole filter.*

The filter in ϕigure _a *has a high input impedance and can be driven from any source; however, it shows more high-frequency feedthrough—about –65 dB—than the other two configurations. The filters in* ϕigures _b *and* _c *have less feedthrough but are sensitive to the impedance of the signal source. This sensitivity alters the gain of the filter in* ϕigure _b *and the frequency response of the filter in* ϕigure _c, *so those two filters should be driven from a low-impedance source.*

All three filters have resistors that you can use to set the filter's Q. Make sure that, with the expected input, you don't exceed the multiplier's peak output current or the amplifier's peak voltage-output limits. The first multiplier is exposed to any large out-of-band signals that may be present. Choose component values so that the multiplier always remains in its linear region.

Typically, the resistors are approximately 1000Ɣ, and the capacitors are 20 to 100 pF on the highest frequency range. You can switch in capacitors as large as several microfarads to achieve a lower frequency range. Switched-range versions of these configurations with cutoff frequencies tunable to 20 MHz find use in commercial products. A single-range filter that uses EL2082 multipliers, grounded integrating capacitors, and buffer amplifiers has been tested as high as 50 MHz.

The EL2082 uses a single-ended, 0 to 2V control input rather than the HA-2547's differential input. The EL2082's current gain is equal to the control voltage; that is, the output current can be as much as twice the input current. If the control signal comes from an offboard source, you may need a unity-gain differential amplifier, such as the INA105 from Burr-Brown (Tucson, AZ), to eliminate board-to-board ground-voltage offsets.

One problem common to all analog multiplier ICs is that their absolute gain is not well-specified from device to device. Manufacturers' specifications give tolerances that are generally ±10%, which means that the filter's cutoff frequency and its Q may not be exactly the value that the manufacturer specifies. When you need greater accuracy, you may need a trimmer resistor or a digitally controlled potentiometer to set each multiplier's gain to the correct value.

One way of setting up such a filter is to open the feedback loops and to connect 1% resistors across both amplifiers. With fixed reference voltages applied to both the filter input and the multiplier-control inputs, adjust the trimmers to give the correct dc voltages at the amplifier outputs. This approach is cumbersome, but, with digitally adjusted potentiometers, you can automate it and make it simpler than using a network analyzer and a technician with a trimming tool.

The multiplier-control-voltage characteristic is not perfectly linear, so the cutoff frequency is not an exact linear function of the control voltage. If this lack of exactness matters, then you can use an extra multiplier in a feedback loop to generate a control voltage with the opposite curvature, which you then apply to the filter. This technique makes the reasonable assumption that all multiplier devices of a given type are nonlinear in the same way.

One application in which the tuning error is unimportant is adaptive filtering. In this case, the filter-control voltages are under real-time computer control and continuously vary to optimize the output data bit-error rate, for example.

Range switching gives further extensions

Regardless of which multiplier you use, whenever you need a tuning range of more than about 20-to-1, you must switch the integrating capacitors. When one end of the capacitor is grounded, it is easy to switch in extra capacitors because the switching strays are negligible factors compared with the capacitors. The easiest device for you to use here is a reed or TO-5 relay, but these relays are bulky, expensive, or both. Your alternative is to use a solid-state switch, but this switch introduces its own problems, such as higher on-resistance and stray capacitance.

A preferred range-switching device is a dual four-way multiplexer with an on-resistance of 40Ɣ, such as the DG409 series available from several manufacturers. Because this switch has an enable input, it can select any of five ranges. The highest range uses a fixed capacitor. When you enable the multiplexer, it then

FIGURE 4

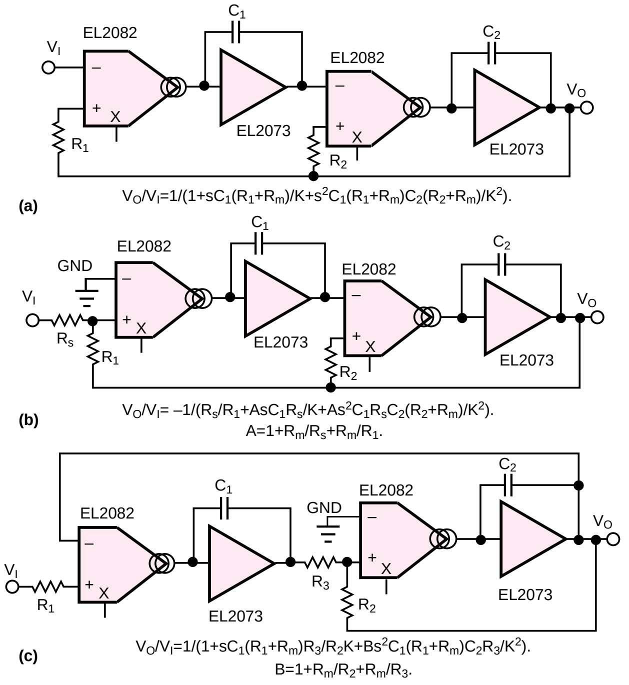

Practical two-pole filters using the Elantec multiplier show some of the possible implementations for this topology. The filter in (a) has high input impedance and can be driven from any source; however, it shows more high-frequency feedthrough than (b) and (c).

switches in one of four parallel capacitors to achieve the lower ranges. One dual multiplexer can switch both of the integrating capacitors in a two-pole filter.

The disadvantage of using an integrating amplifier is that the range capacitors have signals on both ends. The stray capacitance introduced by the switching device and the unused capacitors becomes a problem. Ideally, you would use one multiplexer section at each end of each capacitor.

However, there never seems to be enough board space to do that. Therefore, one multiplexer section has to serve both ends, switching the amplifier input to the various capacitors. The amplifier output now has to drive one end of all the capacitors, whether they are used or unused. The stray capacitance to ground of these inputs may exceed the amplifier's capacitative load specification, and the inputs' stray capacitance to other points in the filter can cause instability. By putting a resistor of about 100Ω in series with the amplifier output, you can reduce these problems without affecting the normal filter performance.

Reaching frequency extremes

Integrating amplifiers are not perfect: They introduce both delay and high-frequency phase shift. It doesn't help that the output capacitance of a multiplier connects to the amplifier input. These effects tend to make the filter have a higher cutoff frequency and a higher Q than you intended. Ultimately, the whole filter oscillates.

Two techniques can extend the filter's useful range. One is for you to introduce a zero into the integrator by adding a resistor in series with the integrating capacitor across each amplifier. Values from 50 to 100Ω significantly extend the integrator's range. The other technique is to connect a capacitor of only a few picofarads from the amplifier's output to the noninverting multiplier input.

When pushing the filter to its limits, you may end up with a design that is normally stable but that oscillates when you turn it on. You can reduce the chances of oscillation by having a multiplier control voltage that rises more slowly than the power supply voltage. Turn-on surges may also overload the amplifier inputs, causing low-frequency, full-scale oscillation unless you have connected clamp diodes to the amplifier inputs.

At the other end of the frequency scale, the filter as a whole has a dc input-output offset on the order of the multiplier's 20-mV input offset. This offset is undesirable in a dc-coupled system. The solution to the offset problem is to connect a difference amplifier between the input and the output to integrate the amplifier's output with a time constant longer than the lowest filter cutoff frequency and then to apply a feedback current to one of the current summing points to balance out any error.

e

References

1. Sallen, RP, and EL Key, "A practical method of designing R.C. active filters," ***I.R.E. Transactions on Circuit Theory****, March 1955.*

2. Russell, Howard T, "Design active filters with less effort," ***Electronic Design****, Jan 7, 1971.*

3. Boyle, Olyn G, "Control an LP-filter's cutoff frequency," ***Electronic Design****, Oct 11, 1976.*

4. Levine, Seymour T, "The rise of active filters," ***Electronic Design****, June 21, 1974.*

5. Napier, Tom, "Multipliers implement tunable filters," ***EDN****, March 16, 1992, pg 131.*

Author's biography

Tom Napier graduated from Aberdeen University (Scotland) with a BS in physics and an MS in electronics. He spent nine years developing spacecraft communications equipment for the signal-recovery group of Aydin Corp (Horsham, PA) and is now a consultant and free-lance writer.