

INCORRECTLY PROCESSED IMAGE-FREQUENCY INFORMATION CAN DISTORT DISPLAYS GENERATED FROM DIGITAL-VIDEO SOURCES. OVERSAMPLING AND WELL-IMPLEMENTED VIDEO-DAC-OUTPUT FILTERS CAN SAVE THE DAY, BUT IMPROPERLY DESIGNED FILTERS CAN MAKE MATTERS WORSE. BEFORE YOU DESIGN YOUR NEXT DIGITAL-VIDEO SYSTEM, TAKE SOME TIME TO INVESTIGATE VIDEO-RECONSTRUCTION-FILTER DESIGN AND TRADE-OFFS IN OVERSAMPLING.

Take the rough edges out of video-filter design

MANY VIDEO APPLICATIONS route video-DAC-output signals through a buffer circuit, an analog-output lowpass filter, or both before sending the signals to a receiving device, such as a TV or a monitor. Before you design your next digital-video application, make sure you understand why you need these circuits in the first place and how to best use oversampling to simplify the filter design and optimize system performance.

The analog output of video DACs requires lowpass filtering to remove unwanted signal components at so-called image frequencies, which digital-to-analog conversion produces. (In this case, the term “image” relates to multiple instances of the signal’s spectrum in the frequency domain and not to images in the pictorial sense.) Removing these frequency components is the reconstruction filters’ main function. These filters prevent aliasing when you connect the DAC outputs to a receiver or monitor that uses an ADC to process the video signal.

In video digital-to-analog conversion, DACs transform 8- or 10-bit digital data into corresponding analog-video signals. The sampling theorem states that, as long as the original analog signal contains no frequencies higher than half the sampling frequency, the signal can be reconstructed completely from

the sampled values. Image-frequency bands appear at multiples of the sampling frequency. To avoid aliasing in the inverse process—analogue-to-digital conversion—you must remove all of the high-frequency components from the signal before you sample it.

Figure 1 shows the frequency spectrum of a video-DAC-output signal. To obtain the signal, the DAC converts to the analog domain an ensemble of digital words. The words are created by sampling at a 13.5-MHz rate, f_{SAMPLING} , a signal whose bandwidth, f_{SIGNAL} , is 6.5 MHz. The lower cutoff frequency of the image-frequency band is $f_{\text{SAMPLING}} + f_{\text{SIGNAL}} =$




Figure 1

A 6.5-MHz-bandwidth baseband signal sampled at 13.5 MHz exhibits this sin x/x behavior.

13.5 MHz–6.5 MHz=7 MHz. The upper cutoff frequency of the image-frequency band is $f_{\text{SAMPLING}} + f_{\text{SIGNAL}} = 13.5 \text{ MHz} + 6.5 \text{ MHz} = 21 \text{ MHz}$. The image-frequency bands repeat at integer multiples of f_{SAMPLING} . Relative to the baseband-signal magnitude, the magnitude of the image bands decreases with each integer multiple, n , of f_{SAMPLING} . This decrease in amplitude is called the $\sin x/x$ effect. At the frequency f_y , where $f_y = (n \times f_{\text{SAMPLING}}) - f_{\text{SIGNAL}}$, you can calculate the signal amplitude as follows:

Signal amplitude $= (\ln x)/x$, where x is expressed in radians $= (f_y/f_{\text{SAMPLING}}) \times \pi$.

For this example with a baseband signal of 6.5 MHz and a sampling frequency of 13.5 MHz, the image magnitudes decrease as follows: –4.25 dB at 7 MHz, –13.59 dB at 20.5 MHz, –17.85 dB at 33.5 MHz, –20.88 dB at 47.5 MHz, –23 dB at 60.5 MHz, and so on.


When the output device does not remove the unwanted frequency components, problems can appear on video signals that contain a lot of fine details. If the receiving device samples at too low a frequency and does not use an antialiasing filter, aliasing can occur and can cause high-frequency, out-of-baseband components to fold back down into the baseband. As a result, unwanted artifacts appear in the video signal.

With a decreasing sample rate, the reconstructed signal becomes increasingly less smooth until, finally, when the sampling frequency becomes less than two times the base-signal bandwidth, erroneous low frequencies appear (Figure 2). Avoiding aliasing requires a minimum sampling frequency of incrementally more than twice the highest frequency that appears with significant amplitude in the signal. Many designers equate “significant” with an amplitude of half of the ADC’s LSB.

REMOVING IMAGE FREQUENCIES


You can use two methods to remove the unwanted image-frequency signals as they emerge from the DAC and to prevent aliasing in the ADC that follows: Use a high-performance external analog low-

Figure 2




This time-domain representation shows how too slowly sampling analog signals creates low-frequency components (aliases) that do not exist in the original data.

Figure 3



Two methods exist for removing unwanted image signals from the DAC output to prevent aliasing in a following ADC: Use a high-performance external analog lowpass filter (a) or digital-interpolation filters and a simple analog lowpass filter (b).

Figure 4



A sixth-order Butterworth lowpass filter with a cutoff frequency of 6.5 MHz adequately filters the two-times oversampled DAC output. The red curve shows the filter’s magnitude response, the blue curve shows the group delay, and the green curve shows the phase response.

pass filter, or use digital-interpolation filters and a simple analog lowpass filter (Figure 3).

High-order lowpass filters are expensive and complex. They also introduce an

unavoidable phase shift onto the signal, causing group delay, which introduces ringing. These filters must be of high order to provide the required attenuation. Because of the reconstruction process’s

$\sin x/x$ effect, the analog filter must also apply peaking to the passband to compensate for the loss in signal amplitude. Obtaining these characteristics further complicates the filter design. Also, the more complex the filter, the greater the likelihood of passband-phase nonlinearity and ripple problems.

Phase shift is an unavoidable characteristic of analog filters. It is the fraction of the period that elapses between corresponding points on the input and output waveforms and is measured in degrees, assuming that 360° represents a full cycle at the frequency of interest. Phase shift indicates how much a signal is delayed. Group delay (in nanoseconds for video filters) indicates how much a filter delays a signal at a specific frequency. Group delay is the derivative of the phase with respect to frequency: $\text{group delay} = d(\text{phase shift})/d(\text{frequency})$.


An ideal filter design would introduce no distortion in the passband. In practice, real filters do not always behave in this way. Generally, as you approach the filter's -3-dB cutoff frequency, f_c , distortion becomes more significant. To avoid this problem, you must increase f_c . Moreover, variations in filter components can introduce additional distortion.

VARYING DELAYS

This behavior can mean, for instance, that a filter with f_c of 7 MHz delays frequencies below 6.5 MHz by a smaller amount than frequencies near 7 MHz. The result is overshoot and ringing at the filter output. High-quality filters exhibit an approximately linear group delay in the passband-frequency range.


In many fourth- to sixth-order Butterworth lowpass filters, the phase stays approximately linear at low frequencies. As the signal frequency approaches f_c , the phase shift becomes nonlinear. Phase shift causes overshoot at the filter output because each frequency experiences a different time delay. A complex filter with delay equalization can deliver satisfactory phase and group delay but is expensive. Such filters require expensive, capacitors and inductors having 1% and

Figure 5



A sixth-order Butterworth lowpass filter with a cutoff frequency of 12 MHz adequately filters the four-times oversampled DAC output. The red curve shows the filter's magnitude response, the blue curve shows the group delay, and the green curve shows the phase response.

Figure 6



Oversampling—in this case, by a factor of two—creates a band of frequencies in which no images exist. Moving the image band away from the signal band substantially relaxes the attenuation-versus-frequency requirements on the lowpass, or reconstruction, filter that follows the video DAC.

lower tolerance. Another method of achieving satisfactory phase and group delay is to add an allpass filter after the lowpass filter. Because an allpass filter does not affect a signal's magnitude but

only its phase, it lets you move the phase inversion to a more convenient frequency (figures 4 and 5).

In video DACs that use interpolation, oversampling can relax the restrictive fil-

TABLE 1—SIXTH-ORDER FILTER SPECIFICATIONS WITH AN f_c OF 7 MHz

Filter type	Attenuation (dB)	Group-delay differences below f_c (nsec)	Phase shift below f_c	Passband ripple (dB)
Butterworth	— 40 at 15 MHz	60	Nonlinear	0
Bessel	— 10 at 7 MHz — 40 at 15 MHz	20	Most linear	0
Chebyshev	— 40 at 11 MHz	270	Highly nonlinear	1
Elliptical	— 40 at 8 MHz	900	Highly nonlinear	1


ter requirements. The principle of interpolation is to insert calculated values into the data stream and thus increase the effective update rate to a multiple of the original data rate. The data stream passes through a digital-interpolation filter with linear phase response, which generates the extra data points. An advantage of finite-impulse-response digital filters is that they can have linear phase response. Interpolation moves the image frequencies to a higher frequency band and thus allows you to use a less complex output filter with a higher cutoff frequency. Another big advantage of using a digital-interpolation filter is that you can more easily compensate for the effect of $\sin x/x$ in the digital-filter response than in an analog-filter design. For example, a digital peaking filter can achieve $\sin x/x$ correction without an interpolator.

When you use interpolation, oversampling also improves SNR because it spreads the noise over a wider frequency band—that is, it extends the width of the noise spectrum to the new, higher, $f_{\text{SAMPLE}}/2$ point. Doubling the sampling rate, also called two-times interpolation, produces a 3-dB SNR improvement. For a 10-bit system $\text{SNR} = (10 \times 6.02 + 1.8) \text{ dB} = (60.2 + 1.8) \text{ dB} = 62 \text{ dB}$. With interpolation, the decrease in total rms noise for a given input bandwidth, f_{INPUT} , is $10 \log(f_{\text{SAMPLING}}/2 \times f_{\text{INPUT}})$. For two-times oversampling, the expression yields $10 \log 27/(2 \times 6.5) = 63.17 \text{ dB}$; for four-times oversampling, the expression yields $10 \log 54/(2 \times 6.5) = 66.18 \text{ dB}$.

SIGNAL BANDWIDTH AND f_c


The video-signal bandwidth changes according to the video standard you use. High-definition video has a different signal bandwidth from standard-definition video, and PAL differs from NTSC video. For standard-definition-video signals, the International Telecommunications Union's ITU-R.BT 601/624 standards provide the bandwidth and data rate as follows. (The letters B, D, G, H, I, K, L, M, and Y following NTSC and PAL denote

Figure 7



The final circuit for the buffer/filter for standard-video output uses an encoder with four-times oversampling and a filter cutoff frequency of 12 MHz.

Figure 8



Oversampling by four times relaxes the reconstruction-filter requirements even further.

national and regional variants of the basic standards.)

- Sampling frequency for NTSC and PAL Y = 13.5 MHz.
- Signal bandwidth for NTSC Y = 4.2 MHz.
- Signal bandwidth for PAL B, G, and H = 5 MHz.
- Signal bandwidth for PAL I = 5.5 MHz.
- Signal bandwidth for PAL D, K, L, M, and Y = 6 MHz.

Designers often work with a signal bandwidth of 6.5 MHz to allow a safety margin for process variations and other errors.

In a video system that uses 10-bit DACs, the SNR, which, in this case, is equivalent to dynamic range, is approximately 60 dB. The video encoder interpolates the data to 27 MHz, giving an oversampling ratio of 2. This conversion rate creates image-frequency bands of unwanted signals from 20.5 to 33.5 MHz,

47.5 to 60.5 MHz, and so on. **Figure 6** demonstrates the requirements for the analog filter; it must have an attenuation of 60 dB at 20.5 MHz if a stopband attenuation of 60 dB is required.

To achieve a 6.5-MHz f_c , you can realize a passive sixth-order filter by using the following components: $C_1 = 180 \text{ pF}$, $C_2 = 680 \text{ pF}$, $C_3 = 470 \text{ pF}$, $L_1 = 2.7 \text{ μH}$, $L_2 = 3.3 \text{ μH}$, and $L_3 = 1 \text{ μH}$ (**Figure 7**). The filter-response plots in **Figure 4** show that the group delay reaches its maximum value, and the phase continues to increase at the higher end of the signal bandwidth. These characteristics introduce overshoot and ringing in the output signal. The attenuation of 59.4 dB at 20 MHz meets the design requirements, however.

To overcome the overshoot and ringing, you can move the filter's cutoff point to a higher frequency, and, consequently, the phase distortion and group-delay peak move to a higher frequency outside

the signal bandwidth. This change requires increasing the filter order, because a sixth-order lowpass filter without stopband zeros provides only -120 dB of attenuation per decade. A simple and cheap analog solution to this problem does not exist. On the other hand, when the video encoder uses four-times oversampling, the situation is different (Figure 8). You can also extend the oversampling ratio to eight times. The ADV7300 video encoder from Analog Devices (www.analogdevices.com) supports eight-times oversampling for standard-definition signals. This rate of oversampling can reduce the lowpass reconstruction filter to a simple third-order Butterworth.


In the four-times-oversampling case, the attenuation requirements are 60 dB at 41 MHz. With these requirements, you can make the cutoff frequency high enough to avoid a group-delay peak and phase reversal in the upper signal band. At the same time, you can provide the required attenuation by using a sixth-order filter without the stopband zeros that typify, for example, elliptical-filter designs. The filter components for this passive, sixth-order, 12-MHz-cutoff lowpass filter (Figure 7) are: $C_1=100$ pF, $C_2=330$ pF, $C_3=270$ pF, $L_1=1.5$ μ H, $L_2=1.8$ μ H, and $L_3=0.56$ μ H. In Figure 5, the phase response and group delay are relatively constant below 6 MHz. The group-delay peak lies well outside the signal bandwidth at approximately 12 MHz. This filter's attenuation is 63.3 dB at 40 MHz.

One other important consideration is the frequency response of the first-stage interpolation filter. The better the filter design, the more frequency components between 6.75 MHz and 13.5 MHz the filter removes. (The number of taps is a good indicator of filter attenuation and passband ripple.) You can find this type of filter in the ADV7300, which uses proprietary SSAF (Super Sub-Aliasing Filter) technology and incorporates $\sin x/x$ compensation.

PASSIVE OR ACTIVE?


Passive filters are designed with inductors, capacitors, and resistors. At high frequencies, passive filters exhibit better behavior than do active filters, whose performance is limited by op-amp bandwidth. Noise from passive filters is en-

Figure 9



When passed through a transmission line, a rectangular pulse with 6-nsec rise and fall times (a) is most faithfully reconstructed when the line is terminated at both ends in its characteristic impedance. Terminating the line at the sending end only increases the overshoot at the rising edge (b).

Figure 10



When you drive a transmission line with a noninverting buffer that has low output impedance, connect a series resistor whose value equals the line's characteristic impedance between the buffer and the line.

tirely the thermal noise of the resistive components, which can be very low. Usually, however, passive-filter input impedance is too low or output impedance is too high, which necessitates the use of additional buffers.

Active filters, which are designed with op amps, resistors, and capacitors, are limited by the op amp's dynamic-performance characteristics, such as magnitude versus frequency. Active filters also require a power supply, which may further complicate the system design. The amplifying stages introduce noise, and cascaded op-amp stages usually make the noise worse, although low-noise amplifiers can minimize noise and allow the addition of gain.

In selecting an op amp, you must consider the finite gain-bandwidth product, distortion, input impedance, differential-gain and phase specifications, and the addition of noise from active components. The dc offset in op amps can also cause problems in active filters. If uncorrect-

ed, dc offset can reduce dynamic range and headroom.

FILTER TYPE

The ideal filter has linear phase response, linear group delay, no passband ripple, and steep rolloff. Different types of analog filters perform well with respect to individual characteristics in this list, but none of the types performs well with respect to all of the characteristics (Table 1).

Chebyshev filters provide steep attenuation despite limited filter order. The main disadvantage of a Chebyshev filter is that the phase is more nonlinear near the cutoff frequency than is the phase of a Butterworth filter. The other disadvantage is that a Chebyshev filter's passband ripple is higher by design than a Butterworth's passband ripple. High passband ripple is the main reason you should avoid using this type of filter in video reconstruction.

A Butterworth filter, on the other

hand, has moderate attenuation for a low-order filter but is still adequate for most video applications. The other advantage of the Butterworth filter over a Chebyshev filter is that the phase is more linear in the passband and cutoff regions. Also, the passband ripple is lower.

A Bessel filter has less attenuation than a Butterworth filter but has highly linear phase in the passband and cutoff regions. The passband ripple is also low. However, to achieve the necessary stopband attenuation without introducing excessive passband attenuation, a Bessel filter's order must be too high. Therefore, Bessel filters are unsuitable for video reconstruction.

The Bessel, Butterworth, and Chebyshev filters belong to the class for which the lowpass function uses an all-pole transfer function; that is, all of the zeroes are at infinity. However, introducing finite-frequency zeros along the real-frequency axis, as in elliptic and Cauer filters, can provide the benefits of a sharper transition from passband to stopband as well

as infinite attenuation at particular frequencies. These improvements are generally achieved at the expense of a reduction of attenuation at high frequencies and an increase in the circuit complexity. The other disadvantage of finite-frequency zeros is poor passband-ripple performance.

Transmission-line interconnections, in which the line's source and load terminations equal the line's characteristic impedance, present a resistive load to the op amp that drives the line. At higher frequencies—that is, at 10 MHz and above—impedance mismatches become increasingly significant. An impedance mismatch between the source and the cable's characteristic impedance causes the op amp's output signal to reflect from the cable's input back into the source. If the load is mismatched, the signal is reflected back from the load into the cable. The result is waveform distortion (**Figure 9**).

If the source impedance is fixed and purely resistive, the load impedance must equal the source impedance to achieve maximum power transfer and minimum

signal reflection. You generally achieve these conditions by matching both the load resistance and the cable's characteristic impedance to the source resistance, which means that the amplifier sees an optimal, purely resistive load and no reflections occur.

This configuration of load and source causes the op amp to deliver half of the input voltage to the load. You can compensate for this attenuation by placing the op amp in a configuration that provides gain. In a noninverting configuration, the op amp provides a gain of 2, which exactly compensates for the voltage division between the source and the load (**Figure 10**).

CIRCUIT DESIGN

When it comes to the actual circuit design, you should pay attention to proper decoupling, grounding, and component selection (**Figure 7**). Route the power-supply leads directly to the amplifier's terminals, ensuring that you keep them to a minimum length. At each terminal, connect a ceramic capacitor

with a value of 0.01 to 0.1 μF in parallel with a high-quality, 3- to 20- μF , tantalum capacitor. Mount the ceramic capacitors as close as possible to the device and the tantalum capacitor as near as possible to the incoming-power pins.

When you choose capacitors for use within filters—that is, for applications other than power-supply bypassing—use devices with the tightest practical tolerances and lowest practical temperature coefficients. Generally, capacitors' value tolerances are 1, 2, 5, and 10%. Also, before selecting capacitors, investigate their high-frequency characteristics. At high frequencies, many capacitors behave in ways that differ from ideal devices.

Resistors for use in filters should be the metal-film type because such devices exhibit less capacitance and stray inductance than do wire-wound types. At frequencies of interest in video filters, most wire-wound resistors behave like inductors or capacitors, exhibiting resonant peaks that make their impedance radically different from their nominal resist-

ance. Resistors should have tolerances of less than 1% and low temperature coefficients.

When choosing inductors, you should consult the specifications for Q (quality factor) versus frequency. The higher the Q, the lower the loss and the more efficient the component. Although you should consider inductors' Q values, most capacitors have very high Q, so, when you choose capacitors, you can often safely ignore this characteristic.

Keep all component leads and etch runs as short as possible. Even though the stray inductance of an inductor's long leads might seem to increase the device's inductance, long leads can be troublesome in inductors used in video filters. The most practical way to control stray inductance is to minimize it. Parasitic circuit elements do not enhance the predictability or stability of any high-frequency circuit's response, and video filters can be especially sensitive.

High-quality video-signal generation relies on preceding the digital-to-analog

conversion with optimal signal processing and on selecting the proper DAC and output filtering. By using digital interpolation and oversampled DACs, you can easily overcome the problems traditionally associated with analog postfiltering and thus achieve the best video performance at reduced system cost. □

AUTHORS' BIOGRAPHIES

Joe Barry is a video-application engineer at Analog Devices (Limerick, Ireland). His responsibilities include managing the video-application and design-evaluation groups and supporting and defining new video products. He holds a BSEE from the University of Greenwich (London).

Christine Bako is an application engineer at Analog Devices (Limerick, Ireland). Her responsibilities include supporting users of audio and video products and writing about those products. She holds a certificate in audio and video engineering from the University of Salford (Salford, England).
