BROADCAST ENGINEERING ■ AUGUST 1997 ■ VIDEO DISK RECORDERS

BROADCAST engineering

An INTERTEC®/K-III Publication

August 1997/$10.00

VIDEO DISK RECORDERS
Can the VDR replace the VTR?

DIGITAL AUDIO
Solutions to routing & mixing

TITLING
Selecting CG technology

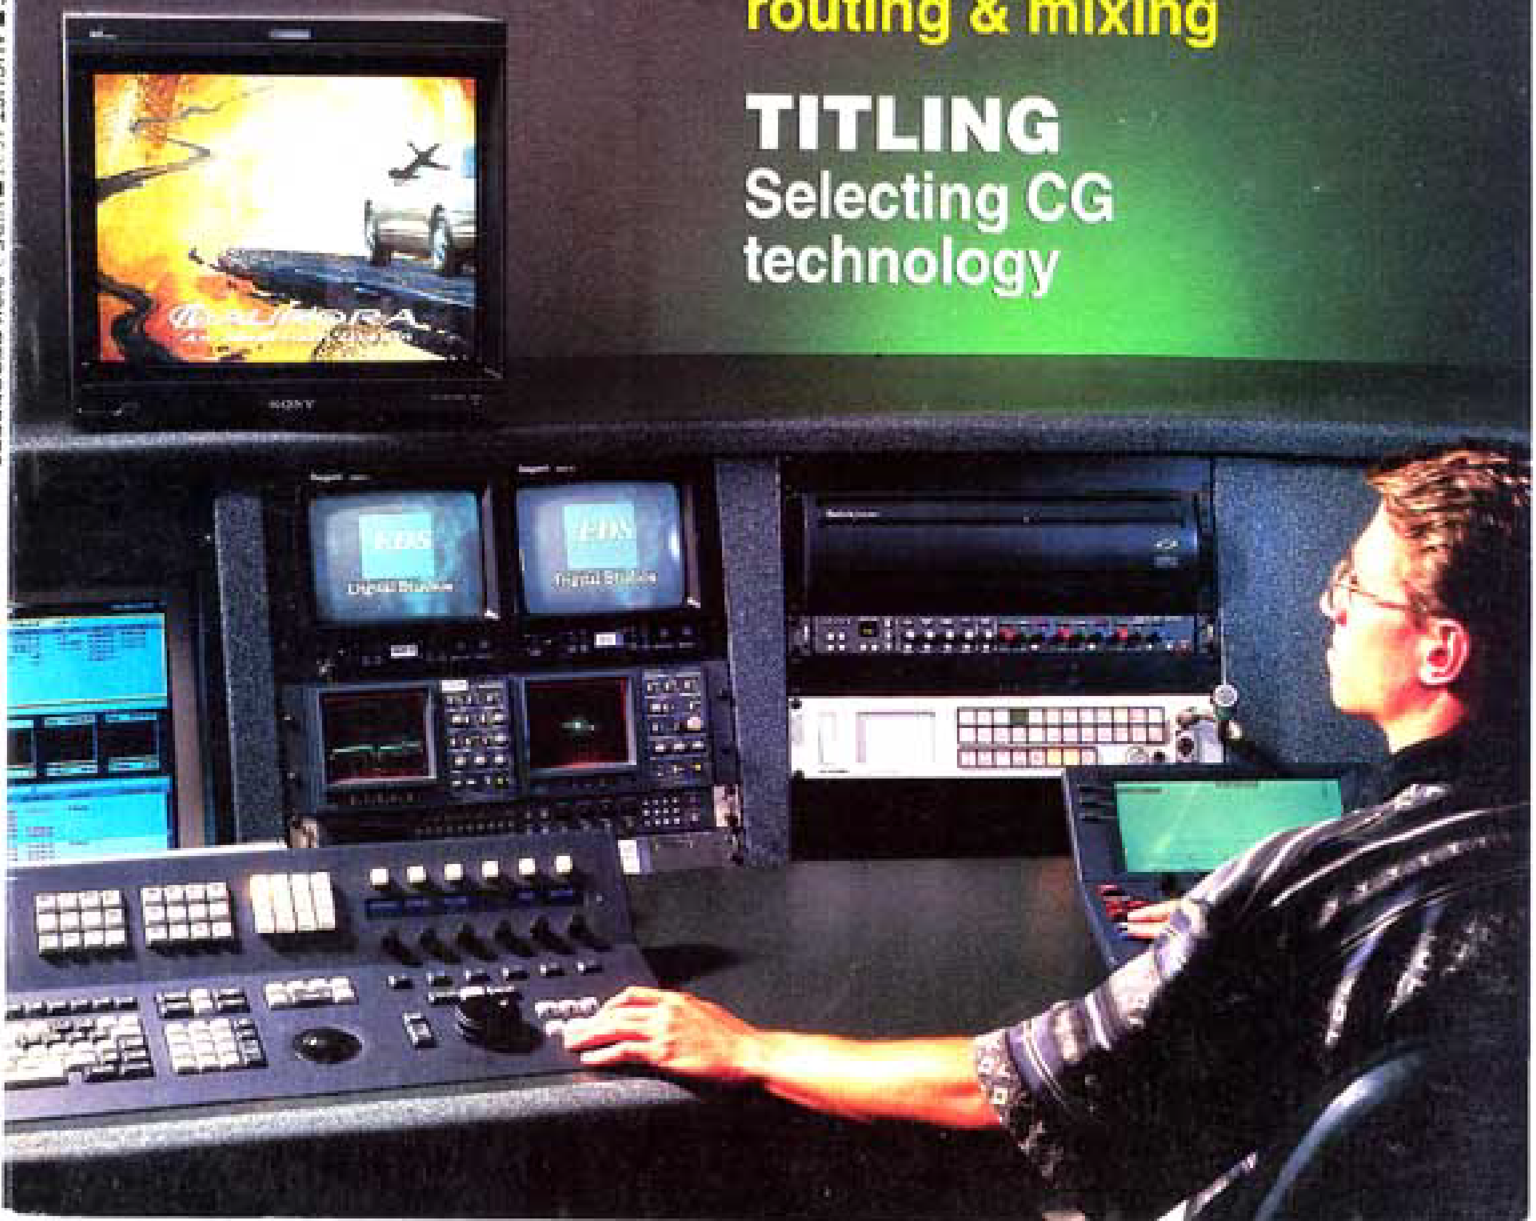

Using DSP technology

Like microprocessors, DSP chipsets are found in many of today's electronic devices.

By Jim Boston

THE BOTTOM LINE:
DSP technology is one reason why today's electronic systems are smaller and lighter than most of their predecessors, despite being more complex. DSP circuits provide powerful signal-conditioning tools in a small package, making it possible to build compact, cost-effective equipment that can be used throughou a wide variety of applications. $

Digital signal processing (DSP) is one of our industry's most overused acronyms. It is generally believed to be a good thing and can be found throughout professional video facilities; digital video switchers, audio mixers, processing amps and digital cameras are a few obvious places. Some not so obvious places include PC modems, some camera robotics systems, cellular phones and VTR video, audio and servo paths. Common DSP functions can be broken down into mixing, filtering, level manipulation and frequency conversion of digital signals. This article looks at why DSP can be beneficial to a system, how DSP performs its magic and how engineers implement a DSP design.

Using DSP technology

Camera DSPs

Before covering the DSP arena in general, let's start by looking at how DSP can benefit a particular piece of equipment, in this case a camera. Figure 1 is a functional block diagram of a DSP camera. It must perform most of the same functions as a non-DSP camera, albeit with a different approach. In older analog tube cameras, the signal was continuous and remained entirely within the analog domain. The pickup tubes developed continuous analog signals from incoming light. Camera electronics processed and combined the signals into a composite signal output. With the introduction of CCD imagers the camera moved from the continuous analog domain into the discrete analog domain. The signal from the CCD imagers was no longer one continuous entity, but a series of samples.

Looking at a CCD's output reveals a series of analog values. These values can change only at a predetermined sample rate. The sample rate, in the case of a CCD camera, is dependent on the number of cells (pixels) in the CCD imager. Today, the number of cells in studio cameras generally ranges from 400k to 640k. The CCD block is basically a sample-and-hold circuit. In the case of the common 520k imager, the array is 1,038 wide by 504 tall (the active picture area is 980 x 494). Photon energy is captured and converted into a volume of electrons representing the amount of light hitting that cell. The resultant 523,152 pixels are marched out 30 times each second, translating to approximately 15.7 million pixels or discrete samples every second. With the addition of horizontal and vertical blanking, the overall pixel rate is ap-

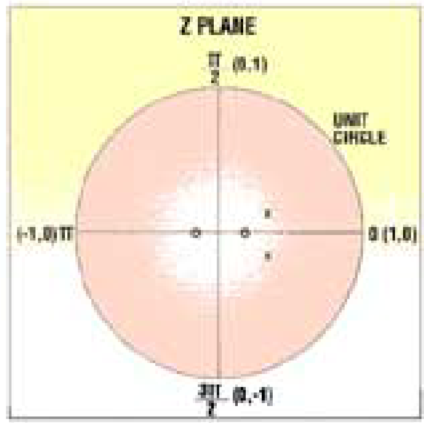

Figure 2. The Z-plane display provides a method to display frequency domain information. Frequencies begin at the zero position (three o'clock) and increase counterclockwise. The zeros and x's represent the zeros and poles (respectively) of a digital filter.

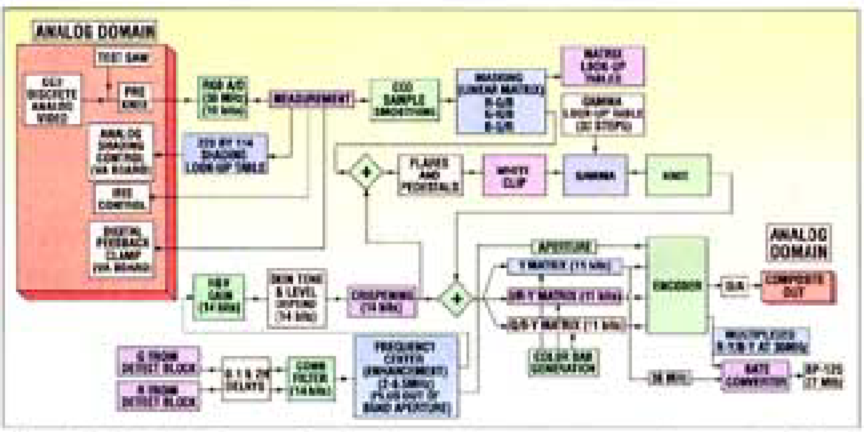

Figure 1. The block diagram of a typical camera shows many areas that can benefit from DSP technology.

Where does analog make sense?

By Kenneth Hunold

Increasingly, digital systems are being installed in facilities across the country (and also around the world) with good reason. Pound for pound (or rack unit for rack unit), digital systems are often smaller and lighter than their analog equivalents. They typically consume less power, are more reliable and are less expensive. Analog facilities are often looked upon as inefficient, substandard and out of date. Is such a characterization valid? Absolutely not! Are there areas where analog systems make perfect sense? Absolutely!

Even if digital transmission systems have significant advantages over analog transmission systems, there are still places where analog systems can be used successfully and appropriately. For audio, if the sources are primarily analog (such as microphones or analog VTRs) and the destinations are analog (such as analog VTRs, satellite uplinks or many fiber circuits), there is no requirement for the production system to be digital. This assumes the production system is reasonably compact and extensive pre- or post-production is not needed. If the material will subsequently undergo extensive post-production, it makes sense to acquire this material in a format (most likely digital) that could later be manipulated with as little generation loss as possible. Nothing here precludes the use of digital controls over analog processes. The concept of digital storage for analog settings using snapshots, pre-sets and scene files allows the benefits of digital control to be applied to analog systems.

For low-level audio and video signals, analog processing is still the only way to go. The tiny signals produced by microphones and video camera pickup devices must be amplified considerably before they can be distributed with comparatively little loss. Microphone pre-amps, even the ones with digital outputs, amplify the signal from the transducer using analog circuits. The types of circuits used still vary, and there are branches of audio philosophy dedicated to the various ways that these analog signals are processed (e.g., vacuum tubes or transistors, discrete devices or integrated circuits, oxygen-free copper or gold conductors, etc.). Low-level video processing, such as black shading and pre-gamma processing, even from such allegedly digital devices as CCD imagers, is still performed in the analog domain. Digital techniques have not progressed far enough to allow these low-level signals (at these frequencies) to be digitized at sufficient resolution to avoid visual distortion. Audio signals are commonly digitized at resolutions approaching 20 bits, while video signals are commonly digitized, for transmission, at eight or 10 bits.

In audio and video effects processing, there appears to be

Continued on page 60

Using DSP technology

proximately 18MHz.

The analog value of the pixel output of these cameras can change every 55.555ns or 980 times per active line. Because this is a discrete signal, it comes with a potential problem — aliasing. When the Nyquist sample theorem is applied to this situation, scenes with the equivalent of 490 vertical lines or more will result in aliasing. To reduce the amount of frequencies above half the sampling rate, the image is filtered optically (and electrically) at the front of the camera. The optical filtering is accomplished through the creative use of birefringence (variable refraction based on light polarization) within pieces of glass. This works out to be an optical low-pass filter.

From the discrete analog domain out of the imager, the signal moves to the discrete digital domain. However, while still in the analog domain, video highlights as high as six times normal are compressed to nearly 200% of normal. This lowers, by two bits, the number needed to handle extreme overexposure. Ten bits can adequately handle the range from 0 to approximately 200IRE. Then each analog signal (R, G, B) is double sampled at 36MHz. The green CCD is spatially offset by half a pixel from red and blue and is delayed by half a pixel when it is converted to a digital signal. Double sampling allows the pixels to be properly aligned spatially. This is one of the "tricks" camera manufacturers use to coax resolutions of more than 800 TV lines from these cameras. (For more information, see "CCD Technology," July 1993.) The analog values are then converted into 10-bit binary numbers (many new cameras are moving to 12-bit samples) and the signal, now digital, consists of three (R, G, B) 10-bit parallel number sequences.

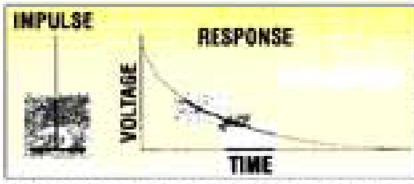

Figure 3. As a design aid, circuitry is tested for impulse response. An analog circuit's response to the impulse shown on the left could be the curve shown on the right.

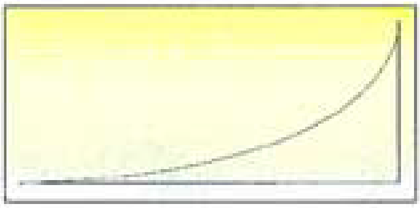
Figure 4. The first step in convolution is to reverse (in time) a circuit's time response. This is the result of reversing the response shown in Figure 3.

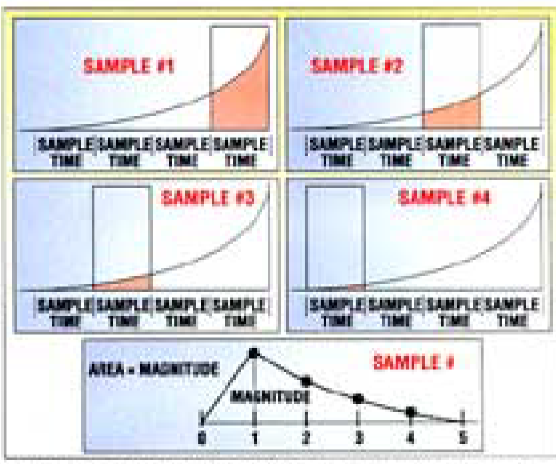

Figure 5. After reversing the time response, the area under the curve is calculated. The results are plotted and form the response curve that an equivalent digital curve must mimic.

Processing in the digital domain

Let's explore how a digital system processes these number sequences. The snap answer is with digital adders and subtractors (or two's compliment addition), right? Actually, the heart of most DSP algorithms is the multiply and accumulate function. This function takes successive digital words (pixel values in the case of a camera) and multiplies them with some value or coefficient at each stage of a DSP pipeline. These values are then passed to the next stage and multiplied with another value. The multiplication results of all the stages are then summed for a final output.

The most common use for DSP

Continued from page 56

a split personality regarding the analog-digital question. Some insist on analog methods to attain the effects that were used in the past and, perhaps, define the concept of "good sound." Others use digital methods to ensure that the effect that was created in a particular session or location can be repeated tomorrow, next week or next year, either in the same facility or in one on the other side of the globe. Some, however, say the effects made 20 to 40 years ago were made in spite of the equipment available then, not because of it — talent always overcomes technology.

At times, it seems analog makes sense only when "digital" has not progressed far enough to be economically implemented. Digital means of modulating AM transmitters have been available for years. FM exciters that operate in the digital domain have become available to complete the digital chain from CD to transmitter. Even the 8-VSB modulation scheme developed for the new digital TV service selected by the ATSC uses analog transmission methods to transmit digital data.

But just as "a rising tide raises all boats," the same rising tide that brings digital techniques to more applications also raises the analog "boat" in the process. Properly designed analog audio circuits can exceed the 105dB (or so) dynamic range of real-world digital circuits. The same advancing technology has also allowed notebook-sized audio mixers to rival the specifications of the best six-figure consoles of a few years ago. (But don't base all of your decisions on the price tag alone.)

Going analog doesn't necessarily mean going cheap. Even the latest high-definition video gear uses analog distribution for reference sync. The tri-level sync signal was developed specifically for analog distribution. Ironically, it might cost extra to integrate this analog signal into an otherwise digital facility, even though it is currently assumed that there is a premium to be paid for "going digital."

Analog or "analogous" signals were developed to mimic real-world phenomenon. Especially in terms of the "transducers" used to capture and radiate sound and light energy, those processes are likely to remain analog for some time to come. However, many of the methods used to process those signals will become digital as larger amounts of processing are needed, and are likely to remain analog when the processing required is minimal. ■

is in filtering. In the lower left quadrant of Figure 1 are instances of bandpass filtering. Additionally, the encoder block must low-pass the two chroma channels to meet output requirements. In the case of cameras that output digital component signals, the data must be converted from a 72MHz parallel byte rate (36MHz for Y, plus 18MHz each for R-Y and B-Y) to the equivalent of RP-125 4:2:2 at 27MHz.

Many DSP designs start with an analog perspective and then transform or equate the analog design to a digital design. For signal processing, it is often easier to work in the frequency domain and some newer DSP ICs can do this. As an example, a DSP system in the frequency domain could demodulate an AM signal simply by breaking a received signal into its frequency components and placing them in "bins." The DSP system would simply have to look at these bins to determine the baseband (demodulated) signal. This is simpler than the envelope detection method used in the time domain.

In the analog domain, functions are usually thought of in terms of frequency. In the early part of the 19th century, Baron Jean Fourier found that overall thermodynamic cycles could be explained by breaking out the individual temperature cycles that combined to make the complex thermal function. The same can be done with electrical functions (i.e., video waveforms). He found that harmonic sine waves of the proper amplitude and phase can be added with a sine wave at the fundamental rate to build any periodic (repeating) function.

A designer of a wideband analog system (like a video camera) might first determine the highest-frequency function a system needs to pass and then create a system to pass the harmonics required to realize a reasonable representation at the output. Designers also typically limit the passband in the system, mainly to minimize noise.

In our simplistic explanation, a transfer function would be determined and a Bode plot (response vs. frequency chart) would be crafted. This is usually done in the frequency domain because in the time domain, the math is much tougher. It is possible to create a mathematical function of the bandpass you want in the time domain, but the problem of creating such a signal or just about any signal, requires capacitance and inductance. Most circuits have both and capacitors tend to integrate the signal over time, while inductors differentiate the signal over time. Differential equations are used to cope with such circumstances, but not many engineers enjoy such endeavors.

There are techniques that convert these problems to algebraic exercises, among them are Laplace transforms. Mathematically, designers can take standard wave shapes in the time domain requiring differential equations and transform them to a simpler algebraic function in the frequency domain. This is done by moving functions that are dependent on time, to ones that are dependent on frequency.

Simply put, a transfer function is a fraction representing the output vs. the input.

$$\text{Transfer function} = \frac{\text{Output}}{\text{Input}}$$

Once a transfer function is reduced to its simplest form, the roots of the numerator and the denominator are found. Roots in the numerator are called zeros and roots in the denominator are called poles. Poles make the response curve in a Bode plot break upward, while zeros make the curve break downward. Generally, the more zeros the faster the rolloff. Realistic transfer functions cannot have more poles than zeros. Once poles and zeros are found, actual values for inductance, capacitance and resistance, along with the system gain can be determined.

This explanation is greatly simplified, and in reality, each block on a block diagram would have its own individual transfer function. In many cases, blocks are broken down into subblocks, each with its own transfer function. Once every block and subblock's function is determined, they could all be multiplied together (as any cascaded system) to determine a system transfer function. However, the result would most likely be too cumbersome to use or understand.

Today, computer tools have replaced the manual analysis required to check designs. A program called Spice takes all the components in a circuit and through nodal analysis (based on Kirchoff's current law) produces time domain and frequency domain analysis of proposed circuit designs. Spice has been in widespread use for more than 20 years and is slowly being replaced by more advanced analysis packages. Today, engineers can describe the external specifications of a proposed design and software can propose the actual circuitry needed. Two of the best-known programs of this type are Verilog and VHDL, both of which are known as hardware-descriptive languages.

A different set of tricks are used for circuitry that does processing in the digital domain. Instead of the Laplace transform, the Z transform is the transform of choice. Like Laplace transforms, the Z transform allows designers to move from the time domain to the frequency domain and vice versa. Instead of a horizontal frequency display (as seen on a spectrum analyzer), the Z transform rolls the spectrum response into a circle, somewhat like a vectorscope display, as shown in Figure 2.

Designers calculate the desired response, craft a transfer function in the frequency domain of the Z transform (see Z transform sidebar) and then perform an inverse Z transform to get back to the time domain. At this point, designers can use the function, now described in terms of time instead of frequency, to find the impulse response. Impulse response is determined by passing a narrow pulse through an equivalent analog circuit. In the case of a low-pass filter the time response might look like the function on the right in Figure 3, with the impulse signal on the left.

A circuit's impulse response can be used to determine the circuit's response to a rectangular pulse. A process known as *convolution* is used to get from one response curve to the other. Convolution works as follows (refer to Figures 4-6):

- First, the time response is reversed in time (Figure 4).
- Next, the square wave is moved

through the response in time steps equal to the sampling rate (Figure 5).

• At each point the area of overlap between the square wave and the response is integrated.

• The area found by integration is used as the magnitude at that point in time.

The digital circuit must mimic this response characteristic. The circuit shown in Figure 6 is one possible solution. The magnitude at each sample point found by convolution becomes a coefficient used by the multipliers. The circuit shown handles only eight bits, but it could easily be expanded. For instance, most professional cameras currently use a minimum of 10 bits and as many as 14 bits to minimize rounding errors in gamma and detail circuits.

Figure 6 depicts a simple digital filter. With only the last four samples summed at any one time this circuit could only mimic a single-pole low-pass filter. To mimic filters with more complex functions, additional samples would have to be summed concurrently, requiring a wider array of registers. These registers are often referred to as taps and some off-the-shelf DSP filters use as many as 55 horizontal taps. Taller arrays are used to handle more bits. Off-the-shelf field-programmable gate arrays (FPGAs) are often used to produce single-chip circuits, such as the one in Figure 6. This allows a single IC type to be used to craft many different filter responses.

The multipliers allow this circuit to act like a low-pass filter. Going back to Figure 5 and examining the resultant response obtained from convolution, the first sample is the highest. If that is called unity (or 1), each succeeding sample is something less than unity. The second is two-thirds of the first, the third is two-thirds of the second and so on. To mimic the required response, multiplier 1 is set to 1, multiplier 2 to 0.66, multiplier 3 to 0.44 and multiplier 4 to 0.29. If an impulse or any other signal arrives at the input, it will produce an output that acts like a simple one-pole low-pass filter. That can be extrapolated to mean this digital circuit has a transfer function matching its analog equivalent. Be aware the circuit shown is simplified, because the gates required to implement the Boolean algebra for the multipliers and the summer are not shown.

This filter is known as a finite impulse response (FIR) filter because it responds to a given input for a finite length of time, in this case, exactly four clock (or sample) periods.

Much simpler would be the infinite impulse response (IIR) filter shown in Figure 7. The transfer function found from convolution had a 0.66 multiplication factor. Therefore, the multiplication shown in Figure 7 could approximate our simple FIR low-pass filter. For Figure 7, the input, output, multiplier, adder and register handle a single bit and must be repeated for the additional bits of an eight-, 10- or even 14-bit signal. The IIR approach works well for simple functions, whereas the FIR approach is needed for more complex functions.

Other DSP uses

Another function necessary in digital cameras is rate conversion, which also uses low-pass filtering. The camera shown in Figure 1 converts a 72MB/s parallel datastream down to the 27MHz rate needed by the component digital output. It would appear all that is needed is to subsample (or throw away some current samples — a process known as decimation) the incoming stream at a three-eights rate and the result would be

Using DSP technology

a parallel bitstream at the correct rate. Interpolation could be used, taking every eight samples, weighing them and creating three samples. But these solutions are no different than sampling the 18MHz pixel rate at 27MHz, which is below the required Nyquist rate and will result in aliasing. Whether the signal is initially undersampled or sampled adequately initially and subsampled later, the results are the same. Rate conversion must use digital low-pass filtering to prevent aliasing.

Another operation performed extensively in DSP cameras is the use of look-up tables. Values in the look-up tables are added to video pixel values. Here, multiplication and the overall summation functions are typically not needed. Areas that make use of look-up tables include:

- *Shading*. Values are added to pixel values based on the pixel's spatial location.
- *Masking*. Correction values are added to each channel's (R, G, B) pixel value to correct for imperfections in the prism's ability to separate the primary colors.
- *Flares*. Correction values are added based on pixel intensity due to optic imperfections.
- *Gamma*. Values are added based on pixel intensity to achieve the desired intensity transfer function. Multiple gamma look-up tables are generally available.
- *Image capture smoothing*. Correction values are used for pixel imperfections in the CCD block.
- *Knee*. Values are added to achieve the desired intensity non-linearity above the desired pixel intensity.

Most camera manufacturers use proprietary DSP chipsets tailored to perform the described camera functions. These sets, often referred to as DSP engines, are used mainly for speed. Many DSP designs use generic DSP processors from a number of manufacturers, such as Motorola, Analog Devices, TI and IBM. These ICs are specialized microprocessors geared for data (such as CCD pixel data) throughput. In contrast to the dedicated DSP chips, which rely heavily on hardware (with lots of internal microcode) for speed, the generic DSP chips rely on a generic hardware architecture and specialized software written for the specific application. This specialized software was written in assembly language, but today it is often written in high-level languages such as C, using compilers supplied by the DSP chip manufacturer.

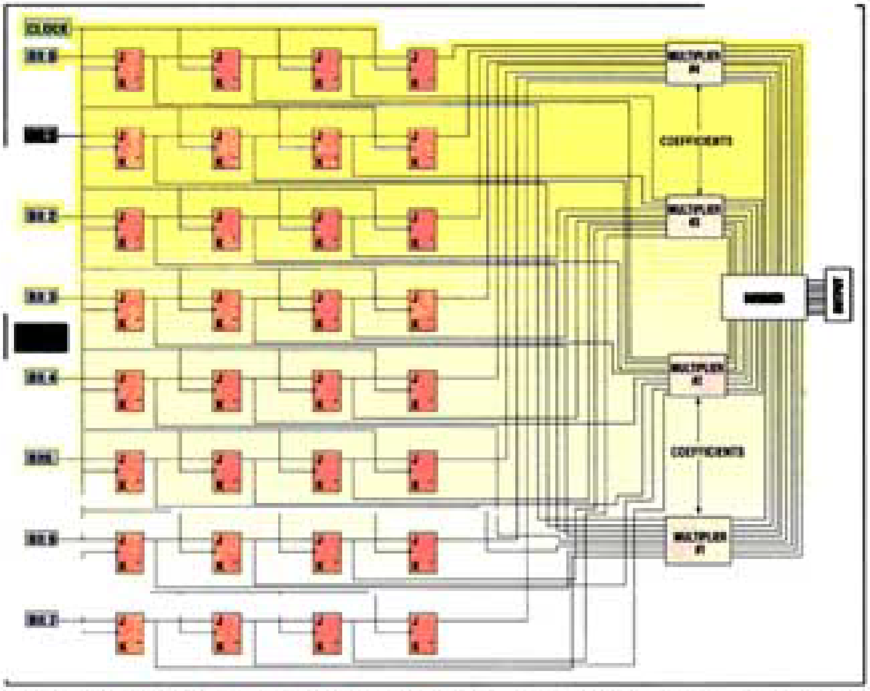

Figure 6. Using digital circuits, a digital equivalent of an analog filter can be constructed. This circuit, with the proper coefficients acts as a low-pass FIR filter.

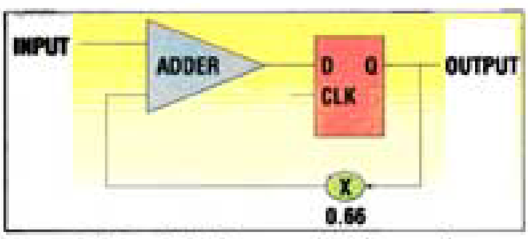

Figure 7. A single-bit low-pass IIR filter can be constructed with a simple set of components.

Implementing DSP designs

Today's generic DSP IC uses an architecture commonly called Harvard architecture, which has separate internal buses for process data and DSP control operations. Some of these DSP ICs use as many as eight separate internal buses. Bus widths can be as wide as 64 bits, but 16 and 32 are the most common. The multiply and accumulate sections of these ICs can have data paths as wide as 96 bits.

To increase throughput, most DSP ICs can perform multiple operations and instruction caches as large as 2k can be found. Although clock speeds can be as high as 200MHz, simultaneous command execution allows a single IC to perform on the order of 1.6 billion instructions per second. Most ICs also allow commands to be issued once and repeated many times. Today, many of the bottlenecks are not due to data processing, but are caused by data movement around the IC or between other DSP ICs that comprise a system. Some implementations leave the data in a common memory and allow other DSP components to look at and change, but not to physically move the data.

The $2.5 billion generic DSP chip industry still has many advantages over what can be done with a standard PC or even a workstation platform. In fact, more than half of today's generic DSPs are used in PCs. The DSP market is expected to quadruple by 2000. More and more, expect to see these ICs in various forms of professional TV equipment. ■

Jim Boston is a consulting engineer based on the West Coast.