DigiPoints

The Digital Knowledge Handbook

Volume III, Issue 3 June 1999

Inside

The Dawn of DSP in Television

DSP Comes of Age in Television

The Technology of DSP Design in Filters

A Glimpse into the Future

For the Reader Who Wants to Know More about DSP Design

Implications for Technical Personnel in Cable Telecommunications

Bibliography

Learning Just Enough to be Dangerous: Glossary

Testing Your Knowledge

Answers to Issue 3-2 Questions

©1999 by the Society of Cable Telecommunications Engineers Inc. All rights reserved. *DigiPoints* is published monthly by the Society of Cable Telecommunications Engineers, 140 Philips Road, Exton, PA 19341-1318, 610-363-6888, June 1999, Volume III, Issue 3. Office of publication is 140 Philips Road, Exton, PA 19341-1318.
Developed by KnowledgeLink Inc.
Editor: Marvin Nelson, BCE, SCTE

Digital Signal Processing in the TV

Editor's Note: This June 1999 issue of DigiPoints*, originally intended for May 1999, directly follows the issue dated April 1999 (i.e., Volume III, Issue 2).*

DigiPoints readers are well aware of digital technology applications in the operator's part of a cable system. This was covered extensively in Volume Two. It should be apparent that for cable telecommunications systems the future is digital. The same observation holds true for customer terminal equipment. Digital Signal Processing (DSP) is rapidly becoming an integral part of all customer terminal equipment. In cable's core business, many of the features that have been added to TV sets over the past decade, such as Picture-in-Picture, are only possible because of DSP. IP telephony is just one example of where DSP technology is necessary for future business opportunities.

This issue will introduce *DigiPoints* readers to DSP by discussing its application in the most basic of the consumer terminals in our industry – the television set. Areas of the TV set where DSP

circuitry is used to create digital filters that make higher-quality and new features a reality will be highlighted. Design considerations, such as gate count, sampling, and digital multiplying, will also be covered as background for understanding other terminal equipment that uses DSP. As these topics are read, the reader should begin to recognize applications of the digital theory that were presented in *DigiPoints, Volume One*.

The Dawn of DSP in Television

Prior to 1985, digital signal processing in televisions was virtually nonexistent. Analog ICs were reaching a fairly high level of integration, and microprocessors were common, but it was hard to justify DSP. There was, however, plenty of research.

DSP's Introduction to Television

Figure 1 shows a block diagram of a high-end television making early use of DSP. It is a simplified representation showing the two areas where DSP was first used. Before talking about the specifics of DSP in the comb or PIP (Picture-in-Picture), system components and their functions will be briefly described.

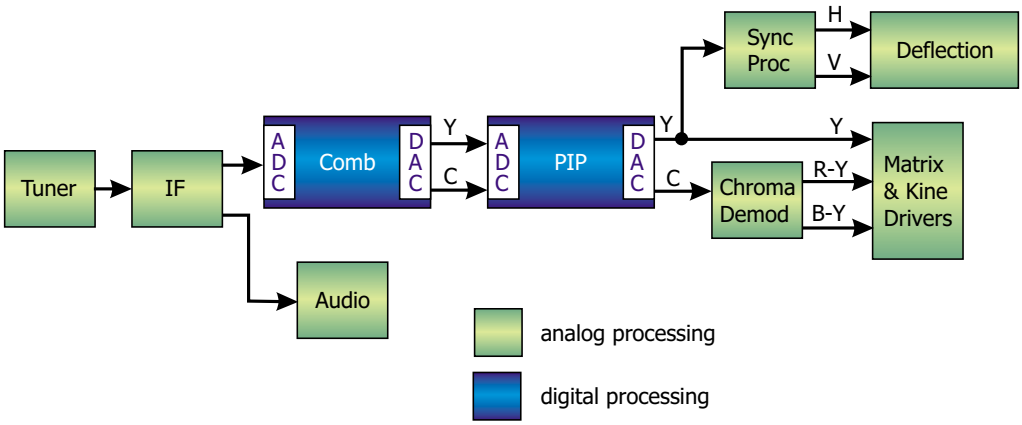

Figure 1: Block Diagram of TV Making Early Use of DSP

The tuner amplifies and down-converts RF signals from cable (or an antenna) to intermediate frequencies, normally 45.75 MHz for the video carrier. The IF section provides baseband composite video and composite audio outputs.

The composite video signal is separated into luminance (Y) and chrominance (C) components by the comb filter. Luminance is the black and white portion of the video signal, and chrominance is the color portion of the video signal. The chrominance portion of the signal uses a 3.58 MHz subcarrier.

The PIP processor allows an inset picture (from a secondary video source, such as a VCR) to be displayed within a main picture. The PIP processor has analog Y and C inputs and analog Y and

C outputs. This makes sense from a manufacturing standpoint, where a particular television chassis might be built with or without a PIP module.

The chrominance output from the PIP block is demodulated to produce R-Y and B-Y signals in the chrominance demodulator block. These signals, along with the luminance output from the PIP block, are used to form the RGB (red, green and blue) signals required to drive their respective guns in the picture tube. Formation of RGB from Y, R-Y and B-Y is called matrixing.

The luminance output of the PIP processor is also provided to the sync processing block, where it is used to generate horizontal and vertical timing signals. These signals are provided to the deflection processor and are used to generate the horizontal and vertical yoke drive signals, as well as to drive the flyback circuitry. The horizontal and vertical yokes are magnetic coils mounted on the neck of the picture tube and are used to sweep the electron beams from side to side and top to bottom. The flyback transformer provides the high voltage (25 KV) required by the picture tube, as well as many of the lower voltage power supplies used in the television.

Comb Filter

One of the very first television circuits to "go digital" was the comb filter. Comb filters provide a high-performance means of separating the luminance and chrominance components of a composite video signal. Before comb filters, low-pass filters were used to extract the luminance component, and band-pass filters were used to extract the 3.58 MHz modulated chrominance signal. Comb filters provide a much sharper picture than low-pass/band-pass systems, and they greatly reduce annoying picture artifacts, such as color fringing on picket fences.

Figure 2 shows a simplified block diagram of a comb filter. This diagram applies to both analog combs and digital combs – in an analog comb, analog circuits are used, and in a digital comb, digital circuits plus analog-to-digital and digital-to-analog converters are used.

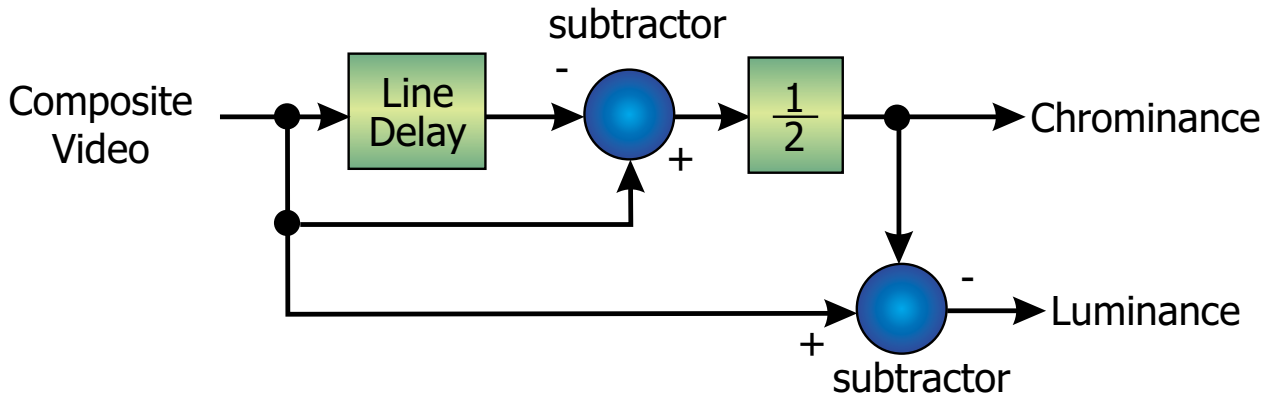

Figure 2: A Simplified Block Diagram of a Comb Filter

© SCTE

The heart of a comb filter is the delay line. The delay is equal to one horizontal line of video, or 63.55 microseconds. The operation of a comb filter is based on the fact that the luminance component of the signal is almost the same on the current line as it was exactly one line ago. Because the chrominance subcarrier has alternating phase from one line to the next with respect to luminance, the chrominance component of the signal is almost the same on the current line as it was exactly one line ago, with the exception that it is inverted.

Figure 2 shows that the chrominance signal is derived by subtracting the previous line's composite video signal from the current line's composite video signal. Subtraction causes the luminance components to cancel out but the chrominance components to double, since subtracting an inverted component is like adding. The luminance signal is then derived by subtracting the chrominance signal from the current line's composite video signal.

Before DSP

The first comb filters were originally analog combs, most using glass delay lines. Glass delay lines are composed of a piece of glass with transducers on both ends. An electrical signal is converted to a mechanical signal at the input transducer, delayed through the glass, then converted back to an electrical signal at the output transducer. Unfortunately, glass delay lines are inherently band-pass devices which restrict the useful frequency range of combing. Glass delay lines exhibit spurious reflections (which degrade the performance of the comb), and they are sensitive to mechanical vibration. Glass delay line combs also require a number of alignments to compensate for wide manufacturing tolerances. Nonetheless, analog combs provided (and still provide) superior performance compared to low-pass/band-pass systems.

After DSP

There are a number of good reasons why comb filters were the first television circuit to go digital. Digital delay lines provide high-performance, adjustment-free delay. More importantly, digital combs open the door to more sophisticated adaptive combing algorithms that further improve pictures and reduce picture artifacts. These algorithms would not have been possible without digital signal processing. And because comb filters were only found in high-end sets, the cost premium of a digital comb could be justified.

© SCTE

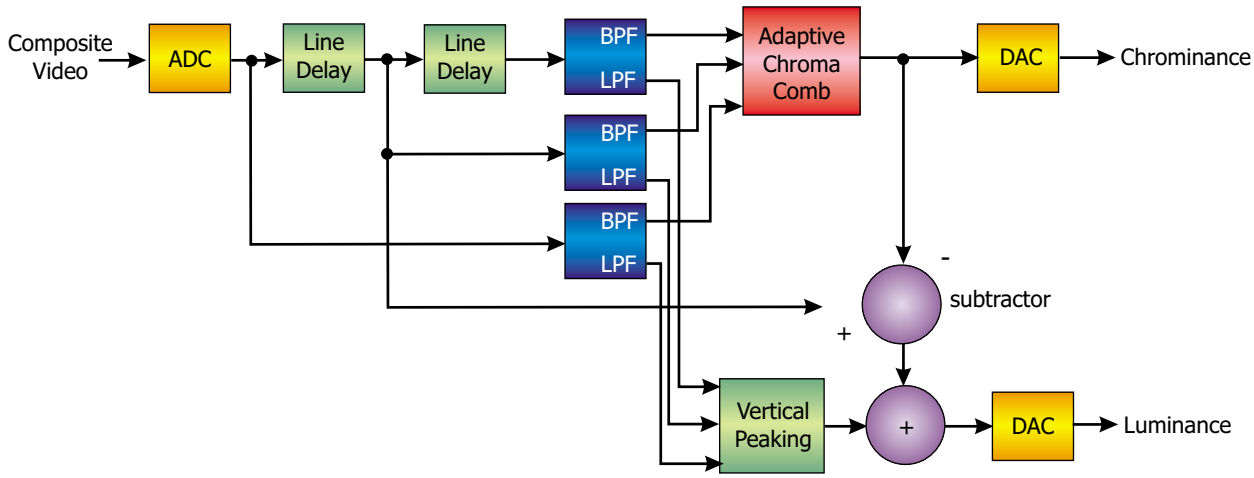

Figure 3: Digital Comb Filter Block Diagram

Figure 3 shows a block diagram of an adaptive digital comb. This comb has an analog-to-digital converter (ADC) at its input and digital-to-analog converters (DAC) at its outputs. This allows the digital comb to mate with the analog circuits in the television.

The digital comb shown in Figure 3 uses two line delays rather than one. This provides three lines of video to work with at a time, rather than two. This makes it possible to adaptively comb in whichever direction is best. For example, if the center line of video most resembles the previous line of video, the center and previous lines will be used for chrominance combing. If the center line of video most resembles the next line of video, the center and next lines will be used for combing. This approach greatly reduces the artifacts associated with one line delay combs.

Luminance is derived by subtracting adaptively combed chrominance from the center composite video line.

Figure 3 also shows another common feature of digital combs — vertical peaking. A vertical peaking signal is derived by comparing low frequency information from the previous, center and next lines of video. This vertical peaking signal is added to luminance to enhance the perceived sharpness of vertical detail in the picture.

Picture-in-Picture

Picture-in-Picture (PIP) was another early use for digital signal processing in television. The reason is simple — unlike other circuits, PIP is only possible with DSP.

The first PIP circuits tended to use DSP for only the bare necessities, such as compression and storage of pixels. Functions like chrominance/luminance separation and demodulation and modulation of chrominance were performed using traditional analog processing. The required

field (or frame) memories were always separate integrated circuits from the PIP processors. Figure 4 shows a block diagram of such a system.

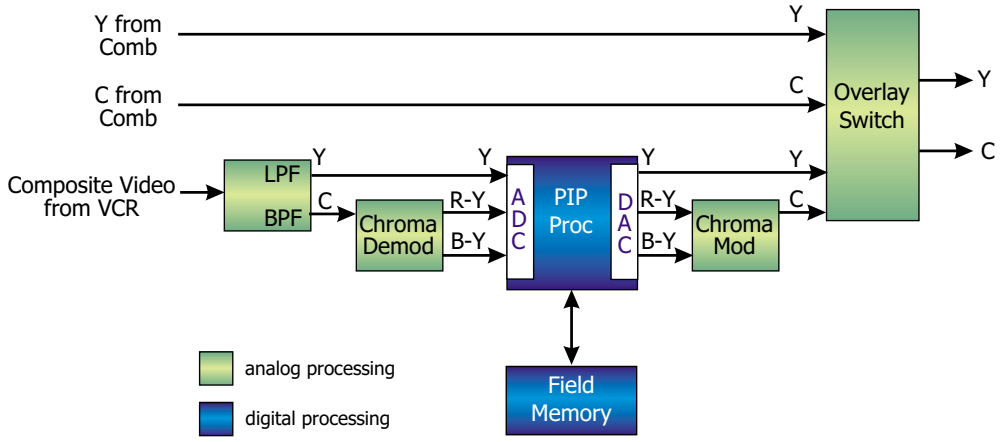

Figure 4: Early PIP Processor

The overlay switch normally selects the Y and C from the comb — the “main” picture path. When it’s time for the scanning electron beams in the picture tube to “paint” the PIP part of the display, the overlay switch selects the PIP Y/C inputs instead of the “main” Y/C inputs.

The main job of the PIP processor is to temporarily store pixels. A field memory is needed to store those pixels. Pixels are stored in the Y/R-Y/B-Y format, which permits easy compression of the PIP picture.

Under worst-case conditions, the very first pixel of a PIP picture may not arrive into the PIP processor until after it’s needed for display. Therefore, up to a whole picture’s worth (a field) of pixels needs to be stored by the PIP processor. The PIP processor will then need to use a stored version of the previous field instead.

Because PIP pixels are stored in the Y/R-Y/B-Y format, Y/C separation and chrominance demodulation are needed to convert the composite video source to Y/R-Y/B-Y. A cheap and simple low-pass/band-pass Y/C separator is considered adequate here. After storage, the R-Y and B-Y outputs of the PIP processor need to be converted back to the 3.58 MHz chrominance format. This is the job of the chrominance modulator.

Even though most early PIP systems relied heavily on analog processing, a notable exception to this was a system designed and used extensively by Thomson Electronics (manufacturer of RCA, ProScan and GE brand consumer electronics), starting in the late 1980s. It performed PIP chrominance/luminance separation, and demodulation and modulation of PIP chrominance, all digitally.

© SCTE

DSP Comes of Age in Television

As the year 2000 draws near, digital signal processing has become the rule rather than the exception. In the 15 years since the introduction of DSP to consumer television, a lot has happened. The price of DSP has tumbled, and digital circuitry has gotten smaller, faster and cheaper. Analog-to-digital converters have come a long way, too. In 1985, it was difficult to get an affordable 8-bit converter to run at 14 MHz. Now, inexpensive 10-bit converters running at 50 MHz are easily accessible.

DSP in Today's Television

Figure 5 shows a block diagram of a modern high-end television and its use of DSP. Before discussing the details of DSP in any of the blocks, the system will be briefly described, including how it differs from the earlier system shown in Figure 1.

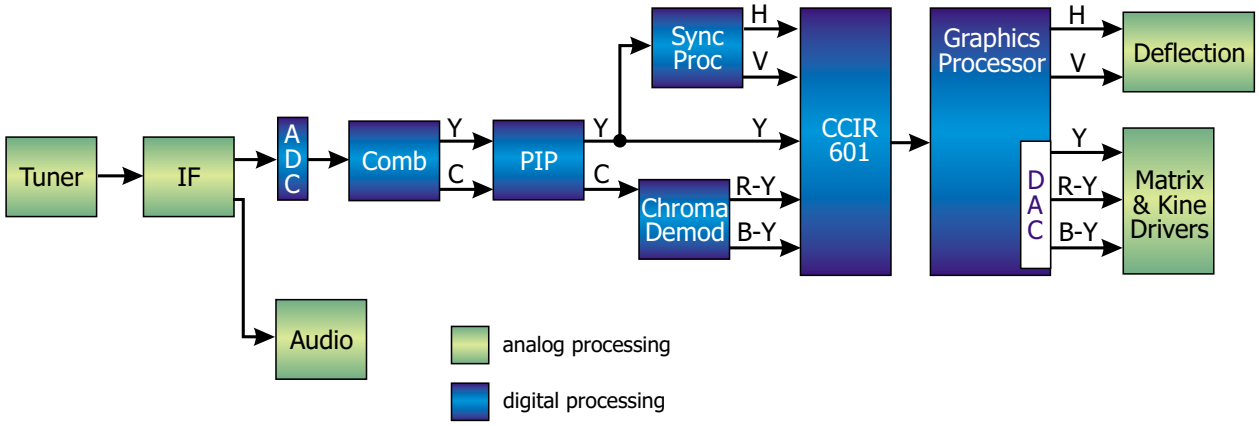

Figure 5: Block Diagram of a Modern TV Using DSP

As in the earlier system, the tuner amplifies and down-converts RF signals from a CATV (or antenna) system to intermediate frequencies (IF), normally 45.75 MHz for the video carrier. The IF provides baseband composite video and composite audio outputs. The composite video output of the IF is digitized, using an ADC. A 10-bit ADC would be typical for this application.

As in the earlier system, the composite video signal is separated into luminance (Y) and chrominance (C) components by a digital comb filter. The digital combs used today are very similar to the early digital combs. In ultra high-end models, so-called "frame combs" using an entire frame of video memory may be used to provide even better performance. A frame comb works similarly to an adaptive line comb, but compares the current composite video line to the

© SCTE

previous line, the next line *and* the line from one frame ago. This approach provides the ultimate comb performance.

PIP processors in modern televisions are also much the same as their predecessors, but much more signal processing is done digitally. The PIP field (or frame) memory is usually integrated into the PIP IC rather than using a separate memory chip.

In the modern television, chrominance demodulation and sync processing are performed digitally. The outputs of the sync processor and chrominance demodulator are provided to a graphics processor via a CCIR 601 interface. The CCIR 601 interface converts the data to the industry standard format for digitized video. This format is used at the graphics processor's input.

The graphics processor is a microprocessor-based graphics generator plus digital signal processor used to create advanced on-screen channel guides. As the number of channels offered to viewers expands, advanced channel guides are becoming a necessity. The graphics processor contains digital-to-analog converters and provides the signals used by traditional deflection and matrix/kine driver blocks.

Chrominance Decoding

As mentioned earlier, chrominance decoding is required to convert the modulated 3.58 MHz chrominance subcarrier into baseband R-Y and B-Y signals. Digital chrominance decoding is very similar to analog chrominance decoding in that analog functions are replaced with their digital counterparts. Figure 6 shows a block diagram of a digital chrominance decoder.

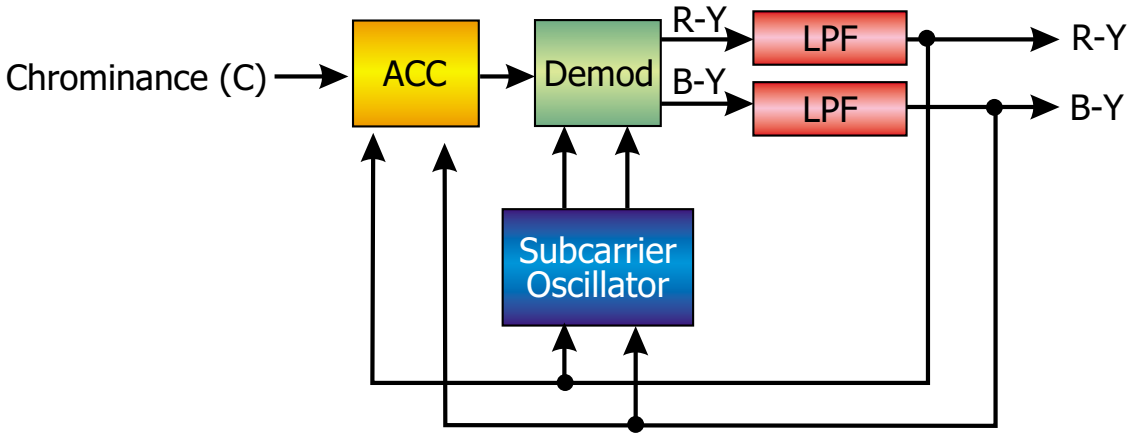

Figure 6: Digital Chrominance Decoder

As in an analog chrominance decoder, the front end of the digital chrominance decoder is the Automatic Color Control (ACC) block. The ACC is an automatic gain control for the chrominance signal. It adjusts its gain to the value needed to cause the chrominance amplitude to meet the nominal value.

The chrominance signal contains a 3.58 MHz amplitude and phase reference as part of the NTSC television standard. This reference is called color burst and is present near the beginning of each line of video. The ACC block looks at the demodulated color burst and adjusts the gain to provide proper amplitude (color saturation).

The heart of the chrominance decoder is the demodulation block. In it, the gain-controlled chrominance is digitally multiplied by a 3.58 MHz sine wave to produce R-Y and by a 3.58 MHz cosine wave to produce B-Y. The sine and cosine waves have a 90-degree phase offset and are phase-locked to the color burst reference. The subcarrier oscillator block looks at the demodulated color burst at the R-Y and B-Y outputs to judge whether its phase needs to be increased or decreased.

This approach mimics an analog system, where a burst-locked chrominance oscillator is used to generate the sine and cosine waves, and where two analog multipliers are used for demodulation.

The outputs of the demodulator block are low-pass filtered to remove unwanted high-frequency components generated as a by-product of the multiplication.

Sync Processing

Embedded in the composite video signal are horizontal and vertical sync pulses. Horizontal sync pulses occur at the beginning of each horizontal line of video and indicate to the television receiver the location of left side of each line. Vertical sync pulses occur at the beginning of each field of video and indicate to the television receiver the location of the top of each field.

The job of the sync processor is to accurately recover the horizontal and vertical timing signals, even in the presence of large amounts of noise. Figure 7 shows a simplified block diagram of a digital sync processor.

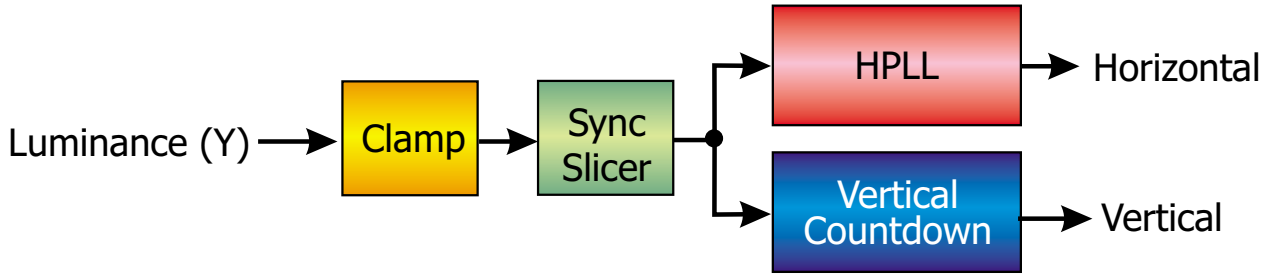

Figure 7: Digital Sync Processor

The sync processor uses the luminance signal from the comb filter as its input. The comb effectively removes chrominance that would make sync recovery more difficult. A clamp is used to adjust the DC offset of the signal to the optimum value for the sync slicer.

The sync slicer is simply a digital comparator. The output of the sync slicer is either a 1 or a 0, depending upon whether the clamped luminance is above or below some predetermined threshold. The output of the sync slicer is called composite sync — it is a 1 during horizontal or vertical sync and 0 otherwise.

Composite sync is sent to both the horizontal phase-locked loop (HPLL) and the vertical countdown block. The HPLL provides a high degree of noise immunity for horizontal sync. Without it, even small amounts of noise would cause pictures to become jagged-looking. The vertical countdown block provides a similar benefit vertically. Without it, common signal degradations would cause the picture to roll.[1]

The Technology of DSP Design in Filters

In analog circuits, signals are in the form of voltages or currents. In digital signal processing, all signals are in the form of numbers. Analog-to-digital converters convert voltage waveforms into a sequence of numbers, and digital-to-analog converters convert sequences of numbers back into voltage waveforms.

In the analog domain, filters, amplifiers and other circuits are built using resistors, capacitors and transistors. In the digital domain, filters, amplifiers and other circuits are built out of digital building blocks such as delay elements, adders, subtractors and multipliers.

Real-World Digital Filters

In the analog domain, filters are built using resistors, capacitors and inductors. In the digital domain, filters are usually built using delay elements and adders. In digital video processing, virtually all filters are of the Finite Impulse Response (FIR) type. (The name finite impulse

1 The output of the HPLL is a square wave, where the rising edge of the square wave is centered in the horizontal sync pulse of video. If there is a missing, noisy or distorted sync pulse, the rising edge of the square wave will be where the center of the horizontal sync pulse should have been. Much more than a noise filter, an HPLL is actually an oscillator that looks at a history of horizontal sync pulses to establish the correct frequency and phase or its square-wave output.

Similarly, the output of the vertical countdown block is a rectangular wave, where the rising edge of the rectangular wave is located at the start of video. If there is a missing, noisy or distorted vertical sync pulse, the rising edge of the rectangular wave will be where the vertical sync pulse should have been. The vertical countdown block works by counting down (thus the name) the number of horizontal lines (262.5) from the last reliable vertical sync pulse.

© SCTE

response just means that there are no feedback loops[2] in the filter.) Figure 8 shows the architecture of an FIR filter.

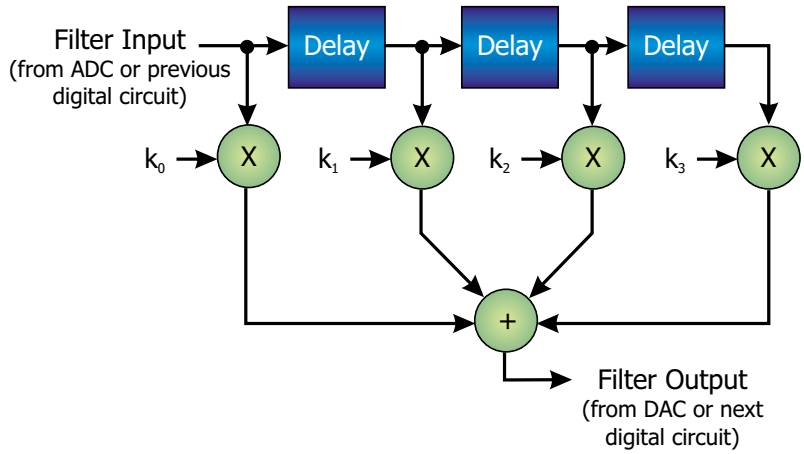

Figure 8: FIR Digital Filter Topology

As shown in Figure 8, a string of digital delay elements is used to create a tapped delay line for the input signal. The delay elements usually have a delay equal to one period of the sampling clock. The filter output is just the sum of the tap values multiplied by their respective coefficients (k_0, k_1, etc.). The value of the coefficients and the number of taps in the delay line determine the filter response. With the right number of taps and the right coefficients, an FIR filter can be used to create any type of filter; low-pass, high-pass, band-pass, notch and equalizer characteristics are all possible. The filter in Figure 8 is a four-tap filter. Some filters have 100 or more taps, but that's unusual.

The main reason FIR filters are used almost exclusively in video processing is that FIR filters with symmetric coefficients automatically provide flat group delay. (Symmetric coefficients mean that the right coefficients are a "mirror image" of the left coefficients. An example of symmetric coefficients for the four-tap filter of Figure 8 would be k_0=5, k_1=9, k_2=9, k_3=5.) Flat group delay is very important in video for two reasons:

- Flat group delay provides a symmetric transient response for luminance and no quadrature distortion for chrominance.
- Flat group delay makes complementary filters possible.

Symmetric transient response for luminance is important for sharp-looking pictures. No quadrature distortion in chrominance is important to prevent color transitions from having objectionable artifacts.

2 Feedback is a design technique that uses an output signal to modify the processing of subsequent input signal variations.

© SCTE

With flat group delay, there is the ability to subtract a low-pass filtered signal from the time-aligned input signal and get a complementary high-pass filtered signal. Complementary filters are often desirable for their performance characteristics, not to mention the fact that they're almost "buy one, get one free." This means you can have a low-pass filter plus a high-pass filter for virtually the same amount of digital circuitry as the low-pass filter alone.

It may come as a surprise, but most FIR filters used in digital video processing are incredibly simple. They're designed to eliminate the need for multipliers and to reduce the number of adders to a bare minimum. For example, the three-tap filter with coefficients 1, 2, 1 requires no multipliers and only two adders. This is because the coefficients are simple powers of two, achievable by circuitry that shifts bits from one position to another. Designers go to great lengths to develop filter characteristics that require no multipliers and a minimum number of adders.

Powers of two aren't the only efficient coefficients for FIR filters. Coefficients that can be built using the sum or difference of powers of two require only one extra adder. For example, a coefficient of 17 would be implemented as 16 times the number plus 1 times the number.

Containing Bit Growth

In digital signal processing, one of the most frequent decisions is how many bits to use to represent a signal at each point in the signal path. If too few bits are used, visible quantization or other artifacts will result. Too many bits waste money because the size of any given circuitry is roughly proportional to the bit width.

The number of bits of a digital signal is analogous to the resolution of a digital voltmeter. On a particular input range, one model of voltmeter might have three digits of resolution. A better meter might have four digits of resolution. In both cases, the range is the same — the difference is resolution.

Once an analog signal is digitized by an ADC, it's normal for it to be processed by a number of DSP stages. Unchecked, each of these stages will probably cause the number of bits to grow. If these excess bits aren't trimmed away, the bit width will grow out of control. To see why signal processing operations naturally increase the number of bits, one may consider the simple digital low-pass filter shown in Figure 9.

© SCTE

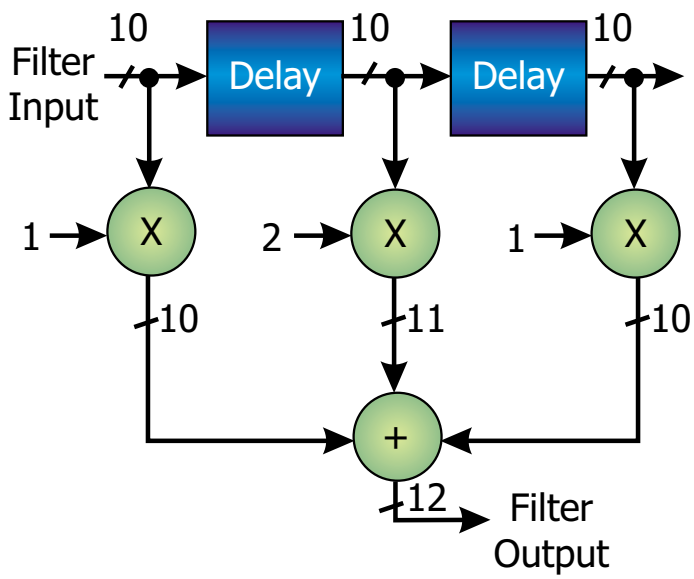

Figure 9: Filter Demonstrating "Bit Creep"

This 1, 2, 1 FIR filter[3] has a 10-bit input signal. Delay elements do not affect the number of bits of a signal, so each of the taps in the delay line still has 10 bits. Now for a look at the outputs of the multipliers. At the first and last multipliers, the coefficients are one. Multiplying by one doesn't change the number of bits, so the outputs of these multipliers are 10 bits. The center multiplier is multiplying by two. This requires an 11-bit number to represent the product.[4]

Looking at the output of the adder, two 10-bit signals and one 11-bit signal are being added. Twelve bits are needed to cover the range of all possible combinations of two 10-bit signals plus one 11-bit signal. This causes the filter output to have two more bits than the input. If this kind of processing were to continue unchecked through a typical video processing IC, a 10-bit signal would quickly grow to 20 or 30 bits!

The solution to this problem is keeping the signal "pruned" by some combination of truncation, symmetric rounding and limiting.

In the above example, all 12 bits at the output of the filter contain valuable information about the signal. Blindly throwing away bits can degrade the signal. But by the time a 10-bit signal has grown to 14 bits through processing, the value of keeping all those extra bits might be questioned.

3 1, 2, 1 comes from the tap coefficients in the filter.
4 It may help to think of how decimal multiplication works. For example, 2 x 5 = 10. The number 5 only requires one digit, but the number 10 requires two.

© SCTE

Truncation, Symmetric Rounding and Limiting

There are a number of ways to reduce the number of bits of a signal. The most common are truncation, symmetric rounding and limiting. Truncation and symmetric rounding are used to eliminate one or more least significant bits (LSBs), while limiting is used to eliminate one or more most significant bits (MSBs).

Truncation simply discards one or more LSBs. This reduces the resolution of the signal and may introduce quantization effects in the picture if enough bits are not retained. A subtle aspect of truncation is that it produces an offset to the signal. In most cases this doesn't cause a problem, but for some signals it does.

A good example is for color difference signals such as R-Y and B-Y. Small DC offsets in these signals produce unacceptable tint shifts for low-level chrominance signals. The alternative to simple truncation in these cases is symmetric rounding. Table 1 shows the difference between truncation and symmetric rounding.

Input	**1-bit Truncation**	**1-bit Symmetric Rounding**
0011 = 3	001 = 1	001 = 1
0010 = 2	001 = 1	001 = 1
0001 = 1	000 = 0	000 = 0
0000 = 0	**000 = 0**	**000 = 0**
1111 = -1	111 = -1	000 = 0
1110 = -2	111 = -1	111 = -1
1101 = -3	110 = -2	111 = -1

Table 1 – Truncation Versus Symmetric Rounding

In Table 1, the zero input row is highlighted for emphasis. First 1-bit truncation will be described. Truncating 1 bit is almost like dividing by two. But notice the point of symmetry has shifted because there are two zero output states, and they are not centered with respect to the input zero state. This is the offset referred to earlier.

The symmetric rounding column has three zero output states that are centered with respect to the input zero state. This produces no DC offset, but does produce a slight nonlinearity about zero. For signals that are "zero-centered," the benefits of zero-offset provided by symmetric rounding generally outweigh any concerns of nonlinearity.

© SCTE

Good examples of zero-centered signals are the R-Y and B-Y color difference signals. When R-Y for a particular pixel is a large positive number, that pixel has a large red component. When R-Y for a particular pixel is a large negative number, that pixel has a large cyan component (cyan is the opposite of red in NTSC color space). Similarly, the B-Y value determines the degree of blueness or yellowness. When R-Y and B-Y for a particular pixel are zero, that pixel should have no color whatsoever. All it takes is a very small offset in R-Y or B-Y to cause pixels that should have no color to appear slightly colored. The human eye is very sensitive to this type of error, so it's important to eliminate offsets in R-Y and B-Y.

Limiting is another way to reduce the number of bits representing a signal. This technique uses specific rules to trim one or more MSBs of a signal, while effectively keeping the signal information intact.

How is Symmetric Rounding Performed?

Symmetric rounding is performed by adding a sign-dependent offset before truncation. Sign dependent means that one value of offset is used when the input is a positive number, and a different value of offset is used when the input is a negative number. For the example in Table 1, a one-bit reduction was shown. This was achieved by adding one to negative numbers and zero to positive numbers prior to truncation.

A Glimpse into the Future

The Future of NTSC

As high-definition television (HDTV) proliferates, existing TV architectures are being replaced with new architectures that bear little resemblance to the old ones. But even with HDTV, it will be many years before today's NTSC signal sources disappear. Consumers will expect to be able to watch their tape libraries using their old VCRs and camcorders for many years. To support this requirement, HDTVs will include an NTSC front-end. This is shown in Figure 10.

© SCTE

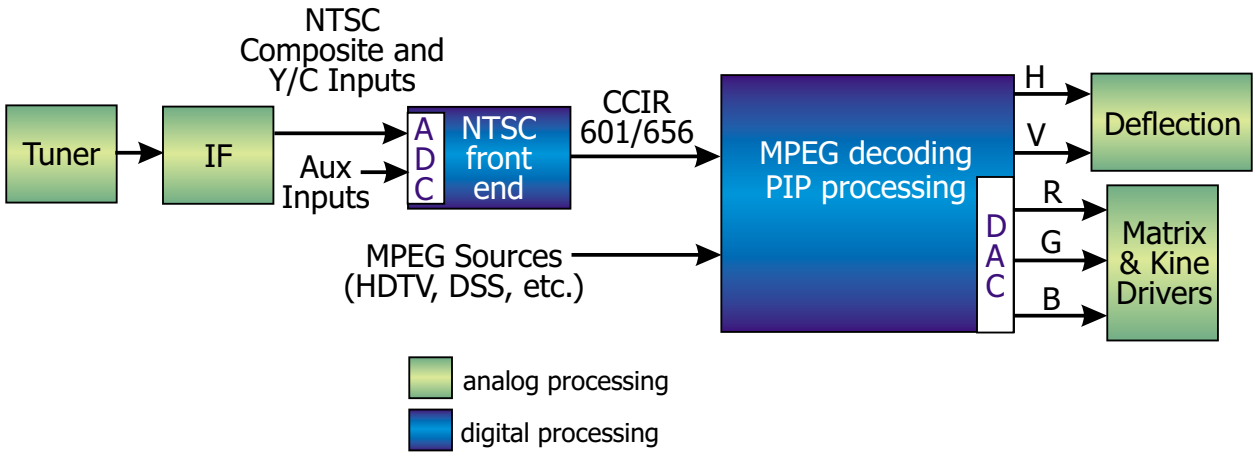

Figure 10 – NTSC Front-End in HDTV

The heart of the system in Figure 10 is the MPEG decoding/PIP processing block. This block receives digital video signals from many sources and in many different formats. It receives MPEG compressed signals from HDTV and DSS (Digital Satellite System) front-ends. It receives digital video from the NTSC front-end using the industry standard CCIR 601/656 interface.

The NTSC front-end receives NTSC composite video and Y/C inputs from an NTSC tuner/IF and from auxiliary input jacks (for VCRs, etc.). The NTSC front-end digitizes its inputs and uses DSP for all signal processing. It includes a comb filter for Y/C separation, a chrominance demodulator and a sync processor. Virtually all the signal processing functionality that goes into today's televisions will find a home in the NTSC front-end.

As long as conventional picture tubes are used for the display, the deflection system and the kine drivers will remain analog circuits. Accordingly, the MPEG decoder/PIP processor includes digital-to-analog converters to provide analog RGB to the kine drivers.

Media Processors

Media processors are software programmable processors dedicated to digital signal processing. They are fairly new, but a lot of big semiconductor manufacturers have announced plans to bring media processors to market.

What media processors have to offer is flexibility. The same media processor used to decode MPEG compressed video when you're watching DSS can be instantly reconfigured to perform NTSC front-end processing when you're watching your VCR. Media processors also offer the potential of faster product development cycles, although in practice this may not be the case.

The true test of whether media processors will find a place in televisions is whether they can do a job more affordably than the alternatives. Time will tell.

For the Reader Who Wants to Know More about DSP Design

Gate Count

Digital designers go to a lot of trouble to keep their circuits as simple as possible. The standard measure of circuit complexity in digital signal processing is called gate count.

A gate is a digital circuit that electronically implements digital logic functions, such as AND, NAND, and OR (see the sidebar). Gate count is the total number of gates in a circuit, where a gate is defined as a standard two-input NAND gate. Other gates can have gate counts higher or lower than this. For example, an inverter might have a gate count of 0.7 gates, and a three-input NAND gate might have a gate count of 1.5 gates. More complex digital building blocks have higher gate counts; for instance, a 12-bit adder has a gate count of about 85. Gate count may also be calculated by counting the total number of transistors in a design and dividing by four.

Minimizing gate count minimizes the physical size of the circuit in silicon, thus minimizing cost. It also minimizes power consumption, which reduces the cost of the power supply and IC packaging. It is common for more expensive IC packages to be required to handle the excessive power dissipation of the circuits inside.

There are a lot of ways for the digital system designer to minimize gate count. Containing bit growth, discussed previously, is one of them. The next few sections explain some of the other basic strategies.

© SCTE

The Dreaded Multiplier

In DSP, multipliers are avoided because they have prohibitively large gate counts. A 12-bit multiplier typically has a gate count of over 1000, whereas a 12-bit adder has a gate count of about 85. That's a factor of 12. And because multipliers are slower than adders, additional gates are often needed for "pipelining" signal processing, or making sure that signals through alternate paths of the circuit arrive at various circuit stages at the same time. This increases the penalty for using multipliers even more.

Digital designers thus try to eliminate multipliers wherever possible. The most common technique is the example that was noted in the real-world digital filters section. The same result as multiplication by a constant that is a power of two (or a sum or difference of numbers that are powers of two) can be achieved economically with adders.

In feedback loops such as automatic gain control (AGC), phase-locked loops (PLLs) or adaptive equalizer circuits, mathematical approximations are often used to eliminate or reduce the need for multipliers. For example, in chrominance processing we often need to calculate:

Magnitude = square root $(x^2 + y^2)$

Digital Logic Functions

The diagrams for commonly used digital gates are shown in Figure A. The tables below show the output of these gates that implement logical AND, NAND, and OR functions. These logic functions are the building blocks of digital circuits.

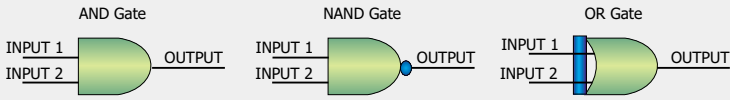

Figure A: Logic Gates

AND Function

Input 1	1	1	0	0
Input 2	0	1	0	1
Output	0	1	0	0

NAND Function

Input 1	1	1	0	0
Input 2	0	1	0	1
Output	1	0	1	1

OR Function

Input 1	1	1	0	0
Input 2	0	1	0	1
Output	1	1	0	1

Table A: Logic Gate Outputs

© SCTE

This would require two multipliers, plus a means of taking the square root. But we can instead use the mathematical approximation:

$$\text{Square root } (x^2 + y^2) \approx 0.5\,(|x| + |y| + \max(|x|, |y|))$$

where |x| is the absolute value of x and |y| is the absolute value of y. This may not look simpler, but this approximation can be performed using two adders and a comparator. (A comparator has the same gate count as an adder.) The point is that we can minimize the number of gates by using mathematical approximations.

Choosing a Sampling Frequency

Remember that the digital representation of an analog signal is derived by the process of sampling, quantizing, and encoding the analog signal. Choosing a sampling frequency may be the single most important decision a digital system designer faces. On the obvious level, the sampling frequency must meet the Nyquist criterion, which says that the sampling frequency must be at least two times the bandwidth of the signal. But there may be good reasons to choose a higher sampling frequency.

Sampling frequency can have a big impact on many digital signal processing tasks. Choosing a sampling frequency of four times the chrominance subcarrier is a popular choice because it greatly simplifies chrominance demodulation, modulation and filtering.

To illustrate the point, Figure 11 shows a chrominance band-pass filter using a 14.318 MHz (four times the chrominance subcarrier) sampling frequency and a similar filter using an 18 MHz sampling frequency. Both filters were designed to meet the same pass-band and stop-band requirements.

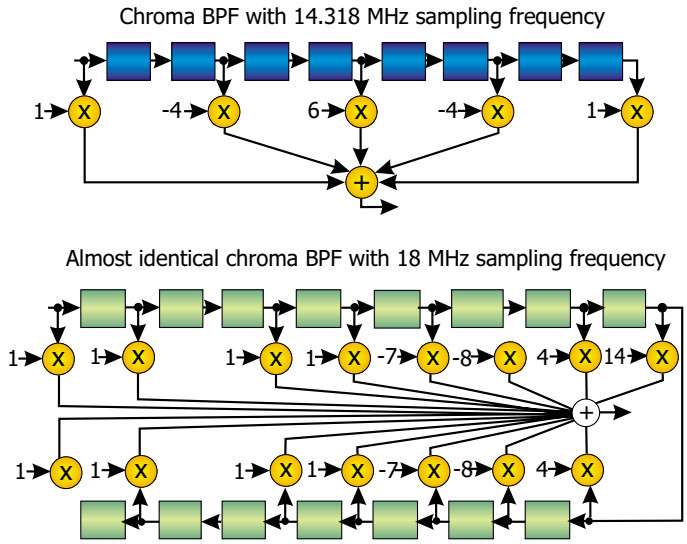

Figure 11: How Sampling Frequency Can Affect Filter Design

Figure 11 shows that the 18 MHz sampling frequency filter requires considerably more taps than the 14.318 MHz sampling frequency filter. This is partially due to the fact that 18 MHz is higher than 14.318 MHz, but is mainly due to the fact that 18 MHz does not have a special relationship with respect to the chrominance subcarrier frequency.

A well-chosen, higher-than-Nyquist sampling frequency can make many digital signal processing jobs a lot easier. But as the sampling frequency goes up, so do the number of gates required; the number of gates is directly proportional to the sampling frequency for many circuits (for example, the line delays in the comb filter). Factors such as electromagnetic interference (EMI) and radio frequency interference (RFI) are also adversely affected by higher sampling rates.

Sample Rate Conversion

Sample rate conversion is converting a data stream from one effective sampling frequency to a different effective sampling frequency. This may seem like an unimportant idea because there is no equivalent operation in the analog (or more correctly, continuous time) domain. But sample rate conversion is one of the most powerful techniques in digital signal processing.

Sample rate converters are becoming one of the fundamental building blocks for digital designers. That is not to say that there is anything like a “one size fits all” design — in fact, every SRC (sample rate converter) must be specifically designed for the application requirements.

Uses for Sample Rate Converters

As pointed out in the previous section, some digital signal processing jobs are easy to do at a specific sampling frequency, but very difficult to do at other sampling frequencies. Good examples are comb filtering, chrominance modulation and demodulation, and sync processing. In cases like these, system designers have the option of using an SRC to convert a signal to an effective sampling frequency convenient for the task at hand. With an SRC, the ADC might operate at one frequency while most of the digital signal processing is performed at a different frequency. Of course, the cost of a SRC often outweighs the cost of performing the task at a sub-optimum sampling frequency.

In most places where sample rate converters are used, there is no alternative. A different sampling frequency is simply required to meet established interface requirements.

How Do Sample Rate Converters Work?

The concept of sample rate conversion is simple. We have samples of a signal at the input sampling rate, but want sample values of that signal at the output sampling rate. This is illustrated in Figure 12. Given the points in black, we need to somehow calculate our best guess for the points in gray.

© SCTE

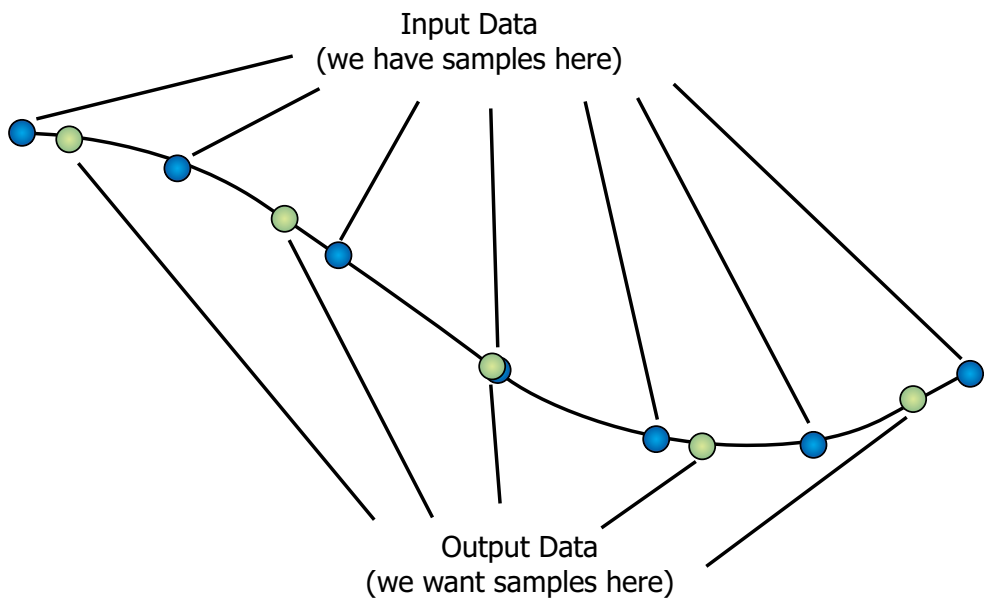

Figure 12: A Job for a Sample Rate Converter

A simple sample rate converter, good enough for some applications, is a linear interpolator. This means that the gray points are calculated by effectively drawing a line between the neighboring black points and finding the appropriate point on the line. The reason linear interpolation isn't perfect can also be seen in the figure. The desired signal at the gray points is not exactly on the line between neighboring black points.

High-quality sample rate converters are built using polyphase FIR filters. These filters use four or more input samples (rather than just the two neighboring samples) to produce an output sample. Polyphase filters have coefficients that dynamically change as a function of where the output point is spatially located with respect to the input points. The filter coefficients are stored in a look-up-table, where they are dynamically selected.

Implications for Technical Personnel in Cable Telecommunications

At first glance, the topics in this issue may appear to be more detailed than needed by the average cable technical person. This may have been the case for our industry before it became telecommunications, rather than television, but things have changed!

Beginning with the television set, we are starting to see radical changes in current and proposed uses for terminal equipment. Existing set-top functionality is migrating toward integration into the TV itself, beginning with simple "cable-ready" sets. The various new formats for digital television signals and displays dictate the addition of yet more functions, just to process a signal for viewing. Like the original set-top box, the devices that provide interfaces between multiple

formats will migrate toward in-set integration. Add the possible interfaces to data applications, and what began as a simple consumer TV will look more like a special purpose computer.

Cable's customers in the 21st century may have purchased and installed any of a number of different equipment configurations just to receive video programming. With the possibility of your customers getting their equipment from retail outlets, rather than from the operator, you will no longer be certain before you enter the customer premises exactly where the TV signal (or other information) becomes the picture on the set. It will help to know some of the underlying technology just to understand why any given device is inserted into the signal path. (Of course, you will also need to understand the difference between the underlying signal formats, such as VSB vs. QAM.)

Increasingly, cable field personnel are going to need to develop analytical skills based on what used to be advanced electronics and digital theory. "Cookbook" solutions to problems will become less common. This applies even more to customer terminals other than the TV set, such as digital telephony and data equipment.

We have used digital signal processing in a state-of-the-art television set as an example of new applications of digital theory that directly affect the service cable telecommunications provides its customers. In many cases, it is impossible to predict the type of other customer terminals you will see in the field over the next five to 10 years. In particular, the material in the section "For the Reader Who Wants to Know More about DSP Design" should help you extend the concepts in this digital signal processing introduction to those future devices.

© SCTE

Bibliography

Jack, Keith. *Video Demystified: A Handbook for the Digital Engineer*. 2nd ed. San Diego: HighText Interactive, 1996.

Robin, Michael and Michel Poulin. *Digital Television Fundamentals: Design and Installation of Video and Audio Systems.* New York: McGraw-Hill, 1998.

Rumreich, Mark. "Pix-in-Pix IC with Digital Decoding/Encoding Using Only One Clock," *IEEE Transactions on Consumer Electronics*, August 1991.

© SCTE

Learning Just Enough to be Dangerous: Glossary

A/D, ADC – Analog-to-digital converter. Video ADCs normally operate at speeds of 13.5 MHz or more.

AGC – Automatic gain control. Used to keep the output level of a circuit constant as the input level varies.

Artifacts – In video, artifacts are noise, spots, streaks or contours induced by the signal processing.

Bandwidth – The range of frequencies a circuit will pass.

Baseband – An audio or video signal that does not have any modulation applied to it.

BPF – Band-pass filter. A circuit that only allows a selected range of frequencies through.

B-Y – In color television, the blue-minus-luminance signal, also called a color difference signal. When added to the luminance (Y) signal, it produces the blue primary signal.

Chrominance – The color part of the of the video signal. In NTSC, chrominance is quadrature modulated on a 3.58 MHz subcarrier.

CCIR 601 – Now known as Recommendation ITU-R BT.601, this is a standard for digitized video. ITU-R BT.601 deals with color space conversion from RGB to YCrCb, the digital filters used for bandwidth limiting, the sample rate and the horizontal resolution.

CCIR 656 – Now known as Recommendation ITU-R BT.656. Interface standard for component digital video.

Comb Filter – A filter that performs Y/C separation. It is called a comb because the frequency response looks like the teeth of a comb.

Composite Video – The combination of luminance, chrominance and timing information in one signal.

D/A, DAC – Digital-to-analog converter.

DSP – Digital signal processing.

Field – An interlaced TV screen is made up of two fields, one containing the even lines, and the other containing the odd lines.

FIR Filter – Finite Impulse Response filter.

© SCTE

Frame – A frame of video is one complete picture out of a video stream. With interlaced video, a frame is defined as two fields, one containing the even lines, and the other containing the odd lines.

Horizontal Sync – The portion of the video signal that tells the television where to place the image on the screen horizontally.

Kine – Short for kinescope – the picture tube.

Luminance – The black and white part of the video signal. Often incorrectly referred to as luminance.

MPEG – Stands for Moving Picture Experts Group. MPEG is a standard for video compression that takes advantage of the redundancy on a frame-to-frame basis of a motion video sequence.

NTSC – Stands for National Television Systems Committee. This is the name of the color television broadcast standard used in North America and some other parts of the world.

Pixel – Short for "picture element." A pixel is a dot or square in a picture.

Quadrature Modulation – The modulation of two carrier components that are 90 degrees apart in phase, by separate modulating functions.

Quantization – The process of converting a continuous analog signal into a set of discrete levels.

RGB – Abbreviation for red, green, blue.

R-Y – In color television, the red-minus-luminance signal, also called a color difference signal. When added to the luminance (Y) signal, it produces the red primary signal.

Saturation – The amount of color present. For example, saturation is the result of the difference between red and pink.

Tint – The color hue. For example, tint is the result of the difference between red and green.

Vertical Sync – The portion of the composite video signal that tells the receiver where the top of the picture is.

© SCTE

Testing Your Knowledge

1. What were the first circuits in televisions to use digital signal processing?

2. What is a comb filter used for, and what did it replace?

3. What is the name of the 3.58 MHz amplitude and phase reference present near the beginning of each line of video?

4. In digital video processing, virtually all filters are of what type?

5. Name three ways to reduce the number of bits of a signal.

6. What are software programmable processors dedicated to digital signal processing called?

7. What is the standard measure of circuit complexity in digital signal processing?

8. What functional element do digital designers try to eliminate or replace wherever possible?

9. What is converting a data stream from one effective sampling frequency to a different effective sampling frequency called?

Answers to Issue 3-2 Questions

1. What are some differences between the NTSC television signal and the European PAL television signal?

   The NTSC signal consists of 525 lines per frame and is scanned at 60 frames per second. The PAL signal consists of 625 lines per frame and is scanned at 50 frames per second.

2. What are the two parts of digital TV technology?

   Digitization of content and viewing format.

3. What is the distinguishing characteristic of Standard Definition TV?

   It has a scan rate of 480 lines per frame.

4. What is the advantage of interlaced scan over progressive scan?

   The same information can be transmitted in one-half the bandwidth.

© SCTE

5. Give two examples of video format conversion.

   Film shot at 24 frames per second can be converted to a 30 frame per second rate by repeating the first of two interlaced fields every other film frame. Scan rates can be converted by using frame stores, which allow the lines from stored frames to be injected between each other, effectively increasing the number of lines per frame.

6. What are the five building blocks of the ATSC digital TV standard?
   - *Video signal format and source coding, which is based on the MPEG-2 compression technology, and a motion-compensated algorithm. The nominal data rate is 19 Mbps.*
   - *Audio signal format and source coding, which uses the Dolby Digital Audio Compression (AC-3) standard, providing 5.1 channels of surroundsound using a 384 Kbps data rate. Multiple audio channels are possible for language and hearing impaired usage.*
   - *Transport and Service Multiplexing, which is the packetized data transport system for combining the video, audio and data. Based on the MPEG-2 standard, packets are 188 bytes long, which includes a 184-byte payload.*
   - *RF/Transmission subsystem, which refers to the channel coding and modulation. The standard ATSC has chosen for the transmission subsystem is 8-VSB (Vestigial Sideband) in a 6 MHz television channel.*

7. What makes VSB different than QAM?

   VSB uses a different form of the carrier signal where only one sideband, and a part (vestige) of the other sideband, is contained in the signal spectrum. Because only one sideband is transmitted, only amplitude, and not phase, can be varied to represent digital symbols. A QAM signal contains both sidebands, and uses both phase and amplitude to represent digital symbols.

© SCTE