Principles of Digital Television Simplified

By E. S. BUSBY, JR.

The principles of signal digitization as they relate to digital television are presented in such a way as to lay groundwork for future digital television papers. To clarify the steps that are involved, a hypothetical "manual" conversion of an analog signal to its digital equivalent is examined. For sampling color television signals, the best sampling rate is considered to be exactly three times the color subcarrier frequency. Various methods of digitization are mentioned, and those suitable for wideband signals are defined. Quantization error and the problems of noise and noise measurement are discussed. The great advantage of handling television signals digitally is that digital quantities are capable of being sent, received, switched, stored, recorded and delayed virtually without distortion. Whatever bad effects accrue from conversion to and from the digital mode need accrue only once. Whatever ills can befall a digital signal, short of total loss, they can be made invisible and they do not accumulate.

Introduction

Digital signals are being used at this time to transmit many thousands of still pictures to ground stations from ERTS (the Earth Resources Technology Satellite). In Europe, experiments are being conducted by EBU members on digital television transmission from satellites. Digital time base correction is already a reality, and digital television signal processing equipment and digital videotape recorders cannot be long in coming. With so much worldwide interest in signal digitization and reconstruction, the engineer who wants to stay up to date in the television field must understand the principles of digitization and have some sense of the probable impact of digitization on the television industry. The purpose of this paper is to discuss digitization principles in such a way that electronics and television engineers will gain a clear if not rigorous understanding that will enable them to more thoroughly profit from future digital television papers.

Doing Signal Digitization "Manually"

To dramatize the process of signal digitization and reconstruction, we might consider how the job could be done "manually" if every step were slowed down 60 million times. On this scale, a microsecond becomes a minute and a wristwatch ticks every five months.

Under such conditions, we could analyze a signal picked up by an antenna by applying the signal to a digital voltmeter (let us say a simple three-digit one with high impedance) through a momentary pushbutton switch, with a capacitor across the input (Fig. 1). When the button is pushed, the capacitor charges to a voltage equal to the signal; between pushes, the voltage on the capacitor is stable and the meter has time to settle down to an accurate reading. It is tedious and inefficient, but we *could* take a reading every 5.6 seconds: push the button, read the meter, write down the voltage — again and again.

Presented on 25 January 1975 at the Society's Technical Conference in San Francisco, by E. Stanley Busby, Jr., Ampex Corp., 401 Broadway, Redwood City, CA 94063. (This paper was first received on 15 April 1975 and in final form on 15 May 1975.)

Finally we could take some graph paper and plot each voltage reading as a point. Then we would draw a smooth curve among the points so that half the points are above the curve and half below it. The result might be a curve such as Fig. 2.

At this stage, let us give names to what we have done. The switch and capacitor performed the function of *sampling,* essentially making the voltage seen during a quick look hold still long enough to be measured accurately. The duration of the button push determines whether we get a sharp or blurred "picture" of that piece of the waveform; the button-pushing rate determines whether we miss capturing any peaks or valleys of the waveform.

The voltmeter performed the task of *analog-to-digital* (A/D) *conversion.* Although the input is analog (in that it can have any value over the operating range), the output is a number that, for a three-digit voltmeter, can have only a thousand values (0–999). If the output reads 0.785 V, the input might have been anywhere between 0.7845 and 0.7855 V. It is in this step that the system accuracy or resolution is forever determined.

The written record is a *memory* or *store.* The numbers could as well have been stored in flip-flops, magnetic cores, or holes in punched cards or paper tape. Plotting points on the graph paper corresponds to making an *digital-to-analog* (D/A) *conversion.* The result is a reconstruction of the original samples, but limited by the accuracy of measurement. Drawing a curve among the points is equivalent to putting the output through a low-pass filter. This curve is as close as we can get to the original signal.

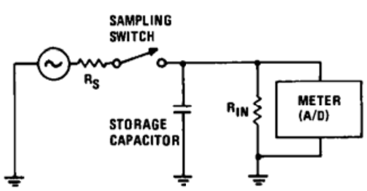

Fig. 1. Essential elements of an A/D converter.

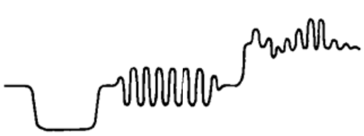

Fig. 2. Reconstructed analog waveform.

Actual Real-Time Signal Digitization

Now let us speed things up 60 million times and get back to "real" time. A real television signal must be sampled at megahertz rates. The voltmeter must settle in tens of nanoseconds. The records are stored in a high-speed electrical memory. The same sequence of sample, wait and record still applies.

The rate that we repeat the sampling, measuring and recording is important. Harry Nyquist showed in 1928 that, assuming a random signal, the original can be reconstructed only if the sampling rate is at least twice the bandwidth of the sampled signal. In the spectrum of a sampled signal, with the sampling rate normalized at two, there is a strong component at the sampling frequency, and like an AM carrier, it has sidebands above and below it (Fig. 3). For an input frequency equal to half the sampling rate, the input and the lower sideband caused by it coincide and cannot be separated. If the ratio of sample rate to sampled input is made greater than two, a low-pass filter can be made to reject the sample frequency and its sidebands. The higher the ratio, the easier the filter design becomes.

In a color television signal, there is a strong component around the color subcarrier frequency. An argument can be made for sampling at an exact multiple, greater than two, of that frequency. Figure 4 shows the sampling rate at three times subcarrier. In the case shown, anything that might go wrong is most apt to do so once each subcarrier cycle, or twice or three times. If twice or three times,

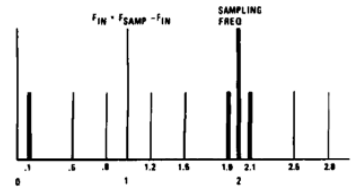

Fig. 3. Spectrum of a sampled signal.

Fig. 4. Sampling at three times the color subcarrier.

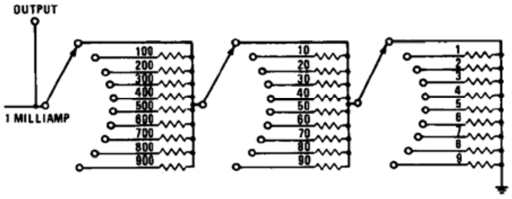

Fig. 5. Principle of simple three-decade D/A converter.

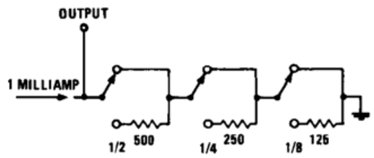

Fig. 6. Three-bit binary D/A converter with 1 part in 8 resolution.

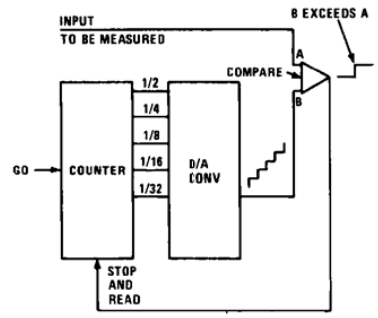

Fig. 7. Counter-type A/D converter.

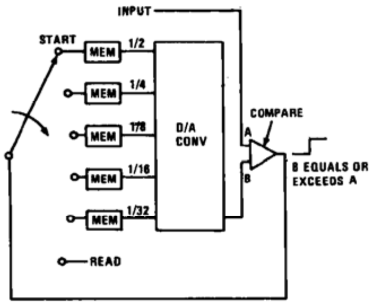

Fig. 8. Successive approximation A/D converter.

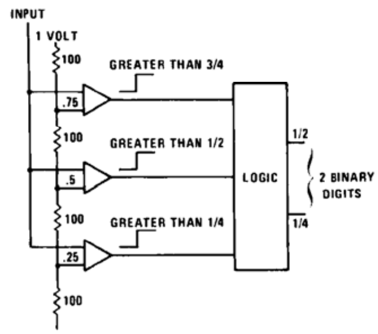

Fig. 9. One-look A/D converter.

the filter will remove the disturbance. If once, at least the disturbance is dot-interlaced and therefore of minimum visibility on the screen.

The color subcarrier in the American NTSC system is quadrature modulated in both amplitude and phase. A short burst of the subcarrier at a reference phase and amplitude is transmitted during horizontal retrace for use in the receiver to recreate a carrier for use in demodulation. A sampling rate which bears a fixed relationship to the reference burst can serve as a continuous reference signal not otherwise available. This can be useful in time-base correctors used on tape playbacks.

In many cases, the signal processing is done digitally and the final numbers are converted back to a signal used to develop a display. This requires digital-to-analog (D/A) conversion, and if we were doing the job "manually" as we did earlier, three 10-position rotary switches connected in series could constitute a hand-operated D/A converter capable of handling three digits (Fig. 5). Numbers in, voltage out.

If we were doing the conversion electronically — say with integrated circuits — it would be more convenient to use binary digits (bits) instead of decimal digits. This can be done by using the integrated-circuit-equivalent of a single-pole double-throw switch to represent each binary digit. Since such a switch has only two possible states, one state can be used to represent a binary "one" and the other a binary "zero." A two-switch, two-bit D/A converter, then, is capable of four levels of output. Adding another switch provides eight levels of output, doubling the resolution (Fig. 6). Four switches yield 16 levels, five switches 32, six switches 64, and so on. Thus, in a typical eight-bit system, the voltmeter reading can be represented by eight binary digits, or bits, to an accuracy of one part in 256.

When any kind of signal is to be processed, the problem of SNR must be considered — and television camera signals, being accompanied by noise, are no exception. A marginal television signal of poor SNR might be adequately processed in a six-digit system because further resolution would only serve to better define the noise and would be a waste. A good quiet signal is adequately handled by an eight-bit system.

Statistically, when the original signal is sampled some samples will land directly on the boundary between one increment of measurement and the next. Such a borderline case can and will go either way . . . a half-increment too high or a half-increment too low. This random choice is measurable as if it were noise. It is sometimes called quantizing noise but is more properly called quantizing error because it need not appear at all unless there is a signal present to cause it. Its magnitude is proportional to the size of the minimum increment. In a binary system, adding one more digit, or bit, halves the minimum increment and reduces the noise measured by six decibels.

Assuming a perfectly quiet input signal of a random nature, there is a rather simple expression for the approximate apparent SNR due to digitization. It is: peak-to-peak signal/RMS noise = −[11 + (6 × No. of bits)], where the p-p signal includes 28.6% sync. In practice, the assumptions have to be adjusted because television signals are not quiet but noisy and they certainly are not random. Because much of the energy of a color television signal is at a high frequency, most of the apparent noise added by digitizing lies outside the video passband. In carefully designed systems, the apparent noise contribution is 8 to 10 dB less than the formula would predict.

Importance of the A/D Conversion

The most critical part of a digital video system is the A/D converter, including the sampler used with it — just as in Fig. 1. The source impedance of the signal must be low enough and the size of the storage capacitor small enough that when the switch is momentarily closed, the capacitor can charge to a voltage insignificantly different from the input. This will require a time several times longer than the product of the resistance and the capacitance (the RC time constant). At the same time, the capacitor must be large enough that the current required to operate the voltmeter during the measurement period doesn't significantly discharge the capacitor.

Of the many sorts of A/D converters, I will describe three basic ones and one which is a combination of two of these. The first sort employs a counter which provides numbers to a *digital-to-analog* converter to produce a voltage which is compared to the input to be measured (Fig. 7). At the beginning, the counter is set to zero. It is then allowed to count until the output of the D/A converter exceeds the input, whereupon the count is stopped and the number is read from the counter. This method is simple, but is much too slow for our purposes. An eight-bit system might have to progress through 255 counts to take a measurement.

The method shown in Fig. 8 is similar in that a D/A converter is used, but dissimilar in that there is individual control over each of the binary digits. At the start, the most significant bit, responsible for a half-scale voltage, is set active, and the resulting voltage subtracted from the input. If the input equals or ex-

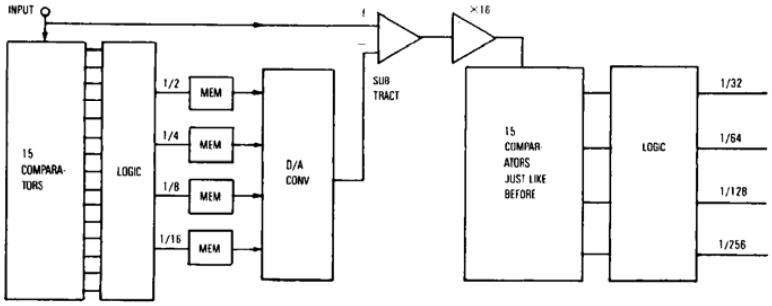

Fig. 10. Two-look, successive approximation A/D converter.

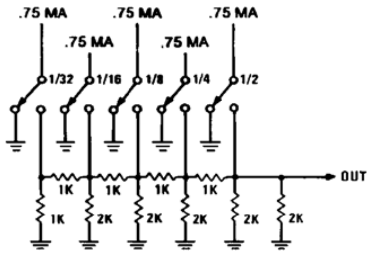

Fig. 12. Improved D/A converter with ladder attenuator.

ceeds one-half scale, the bit is allowed to remain active. If the input is less than one-half scale the bit is set inactive. The next most significant bit, responsible for one-quarter scale, is then set high and the output of the subtractor determines whether *this* bit should be allowed to remain active. After the least significant bit is tested, the number is read from the memories. This method is faster, requiring only eight "looks" for an eight-bit system.

The third basic method (Fig. 9) employs a separate voltage comparator for each possible voltage increment. The number is derived by logic. It is very fast because the number is obtained with only one "look." For an eight-bit system, however, 255 comparators would be required. This method has been used for digitizing video signals, but is cumbersome.

The fourth method (Fig. 10) is a compromise between the second and third basic methods. A 16-level one-look type of A/D converter is employed, yielding a four bit number. These four bits are put into temporary store (memory) and will become the most significant four bits of the final number. In the next step, the four bits drive a D/A converter whose output is subtracted from the input, yielding a difference or error. This error is amplified 16 times and again examined by a similar one-look A/D converter, yielding the least significant four bits. This method is used in a popular digital time-base corrector for videotape playbacks.

Handling the Numeric Information

The amount of numeric information generated by digitizing a television signal is large. For example, an eight-bit system, sampling at three times subcarrier, generates about 86 million bits (86 Mb) per second. Moving these numbers from one place to another can be done bit-by-bit on one wire, much as a teletype signal is sent. Eight wires may be used, one for each bit position, at 10.7 Mb/s on each wire. Or, 24 wires could be used, in three groups of eight, at 3.58 Mb/s on each wire. In computer parlance, the 24-bit group would be called a word, each word representing all the measurements taken on one cycle of sub-

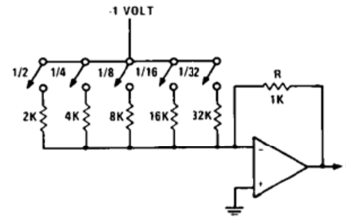

Fig. 11. Simple D/A converter with switches and summing operational amplifier.

carrier. Each eight-bit sub-group of the word is called a byte.

Eventually, we will want to recall our numbers, assemble their bits into the original eight-bit byte and present them to a D/A converter to reconstruct an analog signal. In the first of two popular D/A methods, switches (electronic, *not* electromechanical), each one operated by its associated binary digit, connect appropriate resistors across a voltage supply, producing a current proportional to the size of the number (Fig. 11). The amplifier shown is one way of converting the current to a voltage. Unfortunately, the resistor values cover a wide range, making accuracy difficult to achieve. The other method (Fig. 12) uses identical switched current sources, each connected to its own tap on a ladder attenuator. Each stage of the attenuator divides by two. The least significant digit suffers the most attenuation, the next least significant half that much, and so on. The range of resistor values in the ladder is small, varying only over a two to one range.

At the D/A stage in a system, we begin to see some of the plagues of the analog world. D/A converters are beset with switching transients, ringing, overshoots, crosstalk, etc. Figure 13 shows how a D/A output might vary from the ideal.

While an output filter would remove most of the fast irregularities, some of the discrepancies can show in the filtered output. One method of avoiding these is called "re-sampling." As shown in Fig. 14, the D/A output is examined by a sampler after it has settled down and just before a new number is presented to it. The output of the sampler is then filtered. The filtered result of a sampled or

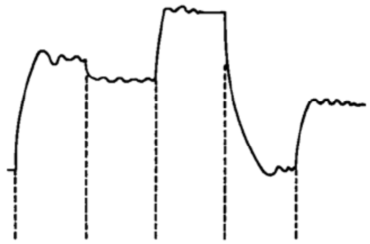

Fig. 13. Nonideal D/A output.

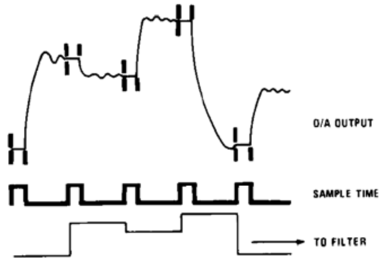

Fig. 14. Principle of resampling.

reconstructed signal exhibits some loss of response at high frequencies that is of the same form as the loss caused by the finite width of the optical soundtrack scanning slit of a sound projector or the finite gap length of a magnetic tape reproducer. The response for any given set of sampled and sampling frequencies is given by sin $\pi x/\pi x$, where x in radians is the frequency being sampled divided by the sampling frequency. In realistic systems, the loss at the highest frequency is only a few decibels and is easily compensated.

Evaluating Digital Television Systems

Some of today's system measurement practices are ill adapted to digital television systems. The practice of measuring noise by shorting out the input and looking at the output does not reveal how the system will behave in the presence of a real signal with real noise on it. Some test signals are of small amplitude and hence are defined by only a small number of digitization levels. The errors encountered are large compared to the test signal and can be interpreted as large system errors, when in fact the errors are insignificant compared to a normal signal amplitude and are invisible on the picture tube. Test signals having no important component less than about 30% of the available signal swing usually produce measurements free of misinterpretation. Unlike typical analog systems, digital systems are at their best when handling full-amplitude signals.

At the cost of bandwidth and some complexity, a digital video system offers in trade some substantial values: whatever is lost in the system need be lost only once . . . in the initial digitization. After that it is all numbers. A binary bit is sturdy. You can't distort it. You may change it completely, or even lose it, but not distort it. You do not adjust a number with a screwdriver.

It is implicit in what we have said that the numbers representing a television signal can proceed through an unlimited number of manipulations without being degraded each time. The computer people have long known that a computer cannot do much with faulty input data; they express this with the acronym "GIGO" standing for "Garbage In, Garbage Out." The converse of this, of course, is "Clean Input, Clean Output" and this is the promise of "Doing it by the numbers."

Digital Television Image Enhancement

By JOHN P. ROSSI

Television image enhancement has long been accomplished by using analog circuits to operate on the analog television signal and thereby emphasize both horizontal and vertical transitions. Now it is feasible to perform comparable image enhancement by digital techniques operating in real time on a pulse-code-modulated NTSC television signal. Several algorithms have been developed to generate vertical and horizontal video details from PCM NTSC television signals encoded at a 10.7-MHz rate (three times the color subcarrier frequency) and at a 14.3-MHz rate (four time the subcarrier frequency). Practical design concepts are presented to implement digital image enhancers. In addition to ensuring the stability, accuracy and reliability that are the usual advantages of digital systems, digital image enhancement will enable video processing techniques to be used that can actually be performed best by digital means. In particular, more nearly optimum coring or "crispening" circuits will be possible and the detail signals will be easily programmable, which may permit automatic adaptive enhancement.

1. INTRODUCTION

Television image enhancement has long been recognized as a means to subjectively improve television pictures. Horizontal enhancement is accomplished by boosting the high frequencies with phase correction for symmetrical overshoots. Vertical enhancement is obtained by the comparison of adjacent lines, low-pass filtering the difference signal and adding the resultant to the main signal. At present, this is done with analog circuits operating on the analog television signal. Recently pictures have been enhanced with digital techniques by computers, but this method is too slow for application to real-time television signals.

Advancements in digital integrated circuit technology now make it feasible to design a digital image enhancer capable of operating in real time on a PCM NTSC television signal. A digital image enhancer costs more than an equivalent analog image enhancer but offers better accuracy, stability and reliability. A digital enhancer also can provide selectable frequency peaking for horizontal details and more nearly optimum coring or "crispening." The present work analyzes digital image enhancement of NTSC signals encoded by pulse code modulation at three times the color subcarrier (10.7 MHz), at four times the color subcarrier (14.3 MHz) and at three times the color subcarrier with phase alternating line encoding (PALE).[1]

Presented, by title only, on 11 November 1974, at the Society's Technical Conference in Toronto; this paper first submitted on 30 September 1974, and in final form on 21 March 1975 by John P. Rossi, CBS Technology Center, High Ridge Rd., Stamford, CT 06905.

2. ENHANCEMENT WITH DIGITAL COMB FILTERING

Image enhancement of television signals by digital methods usually involves comb filtering to generate vertical details and also to remove chrominance components before generating horizontal details. A typical comb filter for NTSC color television signals combines three consecutive (all odd or all even) television lines which we designate T, M and B—for top, middle and bottom—in the following proportions to obtain the chrominance (C) and luminance (Y) signals:

$$C = M - \frac{1}{2}(T + B)$$

$$Y = M + \frac{1}{2}(T + B)$$

A closer scrutiny of this comb filter reveals that the technique is equivalent to sampling and averaging, with certain weighting coefficients, three picture elements from three sequential lines. This sampling and averaging is repeated for all picture elements. A picture element is an infinitesimal image sample that is equivalent mathematically to the Dirac delta function δ. Thus, a pulse-code-modulated (PCM) television signal would seem ideal for comb filtering because each digital codeword describes the instantaneous amplitude of the analog signal at the encode time. The required 1H separation between each of the three video samples is met when the television signal is encoded at a 14.3-MHz rate, or any even multiple of the color subcarrier. This is shown in Fig. 1.

Whenever the signal is encoded at odd multiples of the color subcarrier (e.g. 10.7 MHz), the digital samples in adjacent lines are vertically misaligned, as shown in Fig. 2. Here codeword B_M in the middle line is not in vertical alignment with codewords B_T of the top line and B_B on the bottom line. Comb filtering under these conditions requires the interpolation of codewords B_T, C_T, B_B, and C_B before matrixing them with codeword B_M of the middle line. This interpolation modifies the typical comb filter response in an undesirable manner.

Another method of encoding the PCM NTSC signal at a 10.7-MHz rate is with phase alternating line encoding (PALE). It involves reversing the encoding phase on alternate lines to vertically align the codewords as shown in Fig. 3. The PALE

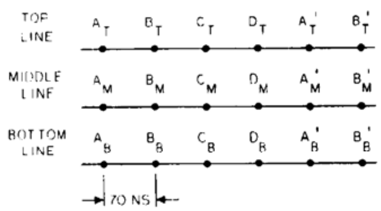

Fig. 1. Digital samples at 14.3-MHz rate from three adjacent scanning lines.

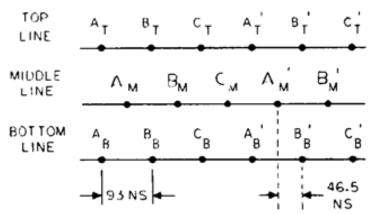

Fig. 2. Digital samples at 10.7-MHz rate from three adjacent scanning lines.

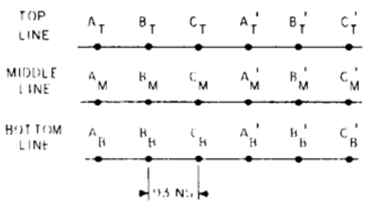

Fig. 3. Digital samples at 10.7-MHz rate PALE from three adjacent scanning lines.