introduced at the extreme left and right edges of the signal. Finally, the amplitudes of the Component 1 luma and chroma signals are reduced slightly, and they are added together to form the Component 1 signal.

Component 2 Encoding

As shown in Figure 4, the side panel signals IS, QS, YSH, and YSD are added and then noise-reduced. These signals are horizontally expanded by a factor of 3.75. Guardbands are introduced at the left and right edges of each expanded panel. The YSH and QS signals are added together and modulated onto one phase of a Fukinuki subcarrier[1] at 3.579545 MHz. The IS and YSD signals are modulated on the quadrature phase of this same subcarrier. The 1 MHz bandwidth of the YSD signal occupies the frequencies above the (narrowband) side panel I signal, and, after time expansion, both fit into the double sideband region and are modulated onto the quadrature axis of the alternate subcarrier. The modulated IS/YSD signal passes through a 3.579545 MHz "inverse Nyquist" highpass filter[2] to allow recovery without crosstalk in the widescreen receiver. Finally, the modulated signals are attenuated and additively combined, forming Component 2.

Shown in Figures 4 and 5 are the insertion gains of the various signals in the ACTV system. The signals on the alternate subcarrier are shown in parallel paths in the lower portions of the picture, identified as "yqs" (for luma and chroma Q component for the side panels) and "yis" (for chroma I component and YSD signals for the side panels). The linear gain factors are shown in the lower left corner of the figures, identified as Q-component side panel gain (QSG), combined side luma and side Q gain (YQSG), I-component side gain (YISG), and an overall encoder gain (EG). These signals on the subcarrier are non-linearly companded both statically and dynamically. In addition, receiver-only processing further aids recovery in noisy conditions.

---

1 A Fukinuki carrier is related to a standard interlaced carrier like the NTSC color subcarrier), but is located in opposite quadrants of the V-T spectrum. Such a carrier can be generated by inverting the phase of an interlaced carrier on alternate fields. Physically, an interlaced carrier appears as an upward moving crosshatch pattern on monochrome receivers, whereas a Fukinuki carrier appears as a downward moving crosshatch pattern.

2 An inverse Nyquist filter has an antisymmetrical magnitude response around a given carrier frequency and is the complement of a Nyquist filter in the decoder. The cascade of the two filters produces a signal that is purely double sideband around the carrier frequency. This property is used to separate the side panel IS/YSD signal from the side panel QS/YSH signal in the decoder.

Component 3 Encoding

As shown in Figure 4, the YHH signal, which initially contains energy in the range 4.5 - 7.16 MHz, is first vertically lowpass filtered to 120 TV lines/picture height (TVL/ph). It is then modulated by a subcarrier at 3.579545 MHz. The modulated signal is followed by a 2.5 MHz lowpass filter to restrict its energy to the range 0.9 - 3.0 MHz. One out of every four lines is time expanded by a factor of 4 so that one line now occupies four active line periods. This time expansion scales the frequency axis by a factor of 4, so that the signal energy now lies in the region 230 - 750 KHz. A final 1.5 MHz horizontal lowpass filter ensures that any remaining repeat spectra caused by the time expansion are fully attenuated. Figure 6 offers a spectral interpretation of Component 3 encoding.

Sync, Burst and Reference Signal Insertion

Normal NTSC sync, blanking, and burst waveforms are added to the signal. Other timing and reference signals used by ACTV are for system timing, ghost reference, and pan-and-scan offset. The timing and ghost reference signals are the same pseudo-random sequence, sent on one line. It is a length 255 pseudo-random sequence repeated three times. The clock rate is 14.32 x 255/256 MHz. The waveform is filtered to a 4.2 MHz bandwidth and sized to have a dynamic range of 100 IRE. Zeros of the sequence are sent as black, ones as white. The sequence is inverted on color field 1.

The reference sequence provides the receiver with an accurate timing reference for raster mapping operations. This sequence allows accurate ghost characterization even in noisy conditions and is designed for computational efficiency in the receiver.

RF QUADRATURE MODULATION AND DEMODULATION

As shown in Figure 2, Components 1 and 2 are additively combined, forming the Main Signal. The Main Signal is RF modulated in quadrature with Component 3, which has been inverse Nyquist filtered. The VSB channel filter limits the Main Signal bandwidth (at baseband) to approximately 4.2 MHz. The filter removes portions of the upper sidebands of the modulated wideband I center panel and the modulated YSH signals as required for transmission.

The quadrature modulated ACTV signal can be received and viewed by a standard NTSC receiver. Component 3 is a signal whose level is quite low for natural images; for ACTV transmission, the video levels of these high frequencies are reduced by 6 dB. This level is so low that receivers with envelope detectors do not show visible interference. If the receiver has a synchronous or quasi-synchronous demodulator, the RF quadrature signal (Component 3) will be automatically removed after passing through the Nyquist filter and demodulator.

A widescreen receiver synchronously demodulates the in-phase and quadrature-phase signals and performs both equalization and ghost cancellation. The recovered Main and Component 3 signals are digitized at 4xFSC and processed by the ACTV decoder, as shown in Figure 2.

Figure 7 shows the RF and V-T spectra of the transmitted ACTV signal. Figure 7 shows how the Component 1 chroma signals (IC and QC) and the Component 2 signals (YSH, YSD, IS and QS) are modulated by the two phase-offset carriers at 3.579545 MHz (FSC = the color subcarrier frequency; FAC = the Fukinuki alternate carrier frequency). Figure 7 shows the V-T spectrum for horizontal frequencies in the range 2.0 - 4.2 MHz. This diagram shows that the modulated Component 1 signals are indeed separate from the modulated Component 2 signals, because the carriers lie in different quadrants in V-T frequency space.

DECODING

Component 1 Decoding

As shown in Figure 5, the demodulated Main Signal passes through a 1.8 MHz bandsplit filter. The highs get intraframe averaged and differenced; the difference signal is modulated Component 2. The intraframe averaged signal is separated into luma and modulated chroma by an adaptive Y/C line comb. The chroma signals are demodulated and lowpass filtered to remove the repeat spectra. The center panel I, Q, Y highs and Y lows are individually compressed by the factor 1.25; these signals are now mapped to their original time slots on the widescreen line. The compressed side panel lows are de-emphasized and expanded by 4 into their original time slots. A lowpass filter follows this expansion to remove and remaining repeat spectra.

Component 2 Decoding

As shown in Figure 5, the quadrature modulated Component 2 signals are processed by two parallel paths. One path contains a 3.579545 MHz lowpass Nyquist filter; the other does not.

The side panel IS/YSD signal is processed *without* a Nyquist filter as follows. First, the signal is demodulated using a quadrature-phase carrier at 3.579545 MHz. The in-phase double-sideband portion of the YSH and QS signal is rejected by this quadrature-phase demodulation; subsequent lowpass filtering is used to reject the single-sideband portion. The demodulated IS/YSD signals are compressed by a factor of 3.75 into their original time slots, and are followed by another lowpass filter to reject the quadrature-phase single-sideband portion of the QS/YSH signal. A 0.7 MHz horizontal bandsplit filter separates IS (the lows) from the YSD signal (the highs). Finally, a gain boost is applied, and the IS/YSD signal is de-emphasized to reduce channel noise.

The YSH and QS signals are recovered in the path containing the Nyquist filter. In this path, the 3.579545 MHz lowpass Nyquist filter converts the IS/YSD component into a pure double-sideband signal, which is then completely rejected by the in-phase demodulator. This Nyquist filter also allows constant-gain demodulation of the QS/YSH signal. After demodulation, the YSH/QS side panels are lowpass filtered to reject the repeat spectra and are compressed by a factor of 3.75 into their original time slots. A 0.7 MHz bandsplit filter separates QS (the lows) from YSH (the highs). Gain boosts are applied, and the YSH and QS side panels are de-emphasized to reduce channel noise.

Component 3 Decoding

As shown in Figure 5, the Component 3 signal, after RF quadrature demodulation, is compressed by a factor of four to create one active line from four. This operation places the horizontal high frequencies into an intermediate band, 0.9 - 3.0 MHz. Vertical interpolation by a factor of four is then used to reconstruct the missing raster lines. A gain boost is applied, and the signal is modulated by a 3.579545 MHz carrier. A final 3.579545 MHz highpass filter is used to remove the lower sideband and recover the band of frequencies, 4.5 - 6.6 MHz. The limiting horizontal frequency, 6.6 MHz, corresponds to a horizontal resolution of 400 TVL/ph.

Decoder Postprocessing

As shown in Figure 5, the recovered side and center panels for each of the I, Q, YL, and YH signals are spliced together in separate splicing modules. Several redundant pixels are transmitted at the side/center boundary; they allow the center and side panels to be "faded" or "spliced" together to reduce the visibility of the seam. Splicing will be covered in detail later. The three intraframe averaged signals (I, Q, and YH) are processed by a time-varying vertical lowpass filter, called a reversing interpolator, to attenuate jagged diagonals. The luma lows, luma highs, and extra horizontal highs (YHH) are added together to reconstruct a luma signal with 400 TVL/ph of horizontal resolution. The recovered Y, I, and Q signals are inverse pan-and-scanned, and each of these signals is then converted to progressive scan. The IFA'd signals are converted to progressive scan using vertical averaging. The luma lows are converted using a more sophisticated motion-adaptive algorithm.

Ghost Cancelling

ACTV's ghost cancelling hardware uses a combination of an FIR filter (for close ghosts) and an IIR filter (for long ghosts). The analysis of the training signal and optimization of the filter coefficients is done in software written, for these tests, in a high-level language. The particular IIR filter delivered for test will cancel up to five discrete "long" ghosts, but the system with its present software will not cancel multiple ghosts if each is of large amplitude. The delay range is from about a 2 microsecond pre-ghost to about a 30 microsecond post-ghost.

Note that the particular parameters of the ghost canceller we deliver for test are essentially unrelated to the ACTV system itself. ACTV recovery is not dependent on ghost cancelling to avoid catastrophic failure - ghost cancelling merely improves the picture, as it does for NTSC.

The training signal is a pseudo-random sequence that requires one line of the vertical interval. The ACTV timing reference uses this same sequence. In the test hardware, the training signal will actually occupy the lowest line of active video.

DETAILED DISCUSSION OF SPECIFIC FEATURES IN THE ACTV SYSTEM

Raster Mapping

Figures 8 and 9 show raster mapping in the ACTV system. The raster mapping operations in the encoder and decoder are *exact inverses* of each other. In Figures 8 and 9, the abscissa indicates active pixel number[3] , where the pixel rate is assumed to be 4xFSC. In the encoder, the center panel signals are expanded by a center panel expansion factor (CPEF) equal to 1.25. The Y side lows are compressed by a factor of 4 into the left and right overscan regions. The Component 2 signals are expanded by a factor of 3.75, as shown in the bottom two waveforms. After raster mapping, guardbands are introduced to insure that the left and right edges of all mapped panels pass through their bandlimited channels with minimal ringing. As shown on the original widescreen signals in Figures 8 and 9, the side and center panels, before mapping, include several overlapped pixels. Because of guardbands, not all overlapped pixels are recovered; however, those that are recovered are spliced together in the decoder to reduce the visibility of a seam.

ACTV delivers a viewable aspect ratio of 16x9. The height of the viewable image is intended to be the same as the height of the viewable image on NTSC receivers, which typical overscan limits to about 450 lines out of the 483 that are transmitted. The width of the viewable ACTV image is set to deliver 16x9 aspect ratio without geometric distortion for an image with a viewable height of 450 lines, consistent with typical NTSC receiver practice.

The resulting 16x9 camera raster, NTSC display raster, and ACTV display raster are shown in Figure 10. The guardbands and taper regions are indicated by the arrows between drawings and by the pixel counts. The central dark gray region is the center panel. The light gray outer regions are side panels. The dashed regions at the outer edges are not transmitted. The intermediate gray pixels between these regions are sent redundantly in compressed form to assure a smooth transition. The indicated horizontal pixel counts are based on a 4xFSC sampling rate

---

3 The numbering system described herein corresponds to the ACTV-I simulations on Sarnoff's DVF. The simulations process arrays 768 pixels wide. Centering the active line within this array allows for symmetrical processing of the signal.

(approximately 14.3 MHz); they represent temporal samples, not necessarily spatial resolution.

A corresponding set of parameters for Component 2 is given in Figure 11. Note that side panel luminance high frequencies and side panel chrominance do not occupy the same number of pixels. Also, the low bandwidth chroma signals require wider taper regions than the luma signals in Figure 10.

Figures 10 and 11 indicate good camera practice for images that must be displayed on NTSC receivers, which requires confining the essential image elements to the height actually displayed, including overscan, on NTSC receivers. Given this height, the "intended viewing area" of the 16x9 image is indicated.

Guardbands

The guardbands for Components 1 and 2 are shown in Figures 12 - 14. These figures show individual pixels for the range of interest, assuming the panning window is centered. Figure 12 shows Component 1 guardbands. The top two waveforms are the expanded luma center panel and compressed luma side panels before guardbands are introduced. The third waveform shows the resulting waveform after guardbands and splicing are introduced. Two different types of operations are used in guardband formation. One operation, called a *taper*, weights original data values near the edge of a panel by a raised cosine-shaped weighting function. The second operation, called a *barrel taper*, augments the edge of the panel by abutting pixels values from the opposite edge of the image. Thus, a taper operation retains the original panel size and merely rolls off the brightness into blanking; a barrel taper operation increases the size of a panel by adding "extra" pixels between the panel edge and blanking. Both of these operations are used to produce Component 1 luma. Barrel tapers are used at each each of the compressed side panels, and a taper is used at the edges of the expanded center panel. The waveforms are added together, effectively producing a cosine splice between the side and center panels. Component 1 chroma gets tapered with a longer weighting function. The shape of the taper function insures that chroma is zero by the time it reaches the monochrome overscan regions; it also insures that the chroma signals will pass through narrow bandlimited channels.

Figure 13 shows guardbands for the Component 2 signals. The top waveform shows the panels after raster mapping. Since the expanded left and right side panels will be, in general, uncorrelated, a smooth transition

must be inserted between the panels. As shown in the middle waveform, this is done by over-mapping the inside edges, and adding these tapered regions together.

Figure 14 summarizes and compares the guardbands for Components 1 and 2 on the same time axis. It is important to note that the lower two waveforms, Component 1 chroma and Component 2, are tapered to blanking *before* they reach the overscan regions. Since the overscan regions are transmitted as luma only, no additional artifacts are introduced into the reconstructed side panel lows.

Decoder Splicing

Figures 15 - 18 show how the various side and center panel signals are spliced together in the decoder. The intraframe averaged signals (I, Q, and Y highs) are spliced separately from the Y lows because the former are processed separately in the unified V-T post-processor.

Figure 15 shows the relevant luma lows signals in the decoder. The top two waveforms show the side and center panel luma lows after raster mapping them into their original widescreen time slots. Both signals are recovered with one or more "bad pixels" at the extreme edges (these pixels were tapered by guardbands in the encoder). Only one "good" overlapped pixel is available for splicing, and this pixel is simply averaged to produce the output pixel.

Figure 16 shows the relevant luma highs signals in the decoder. The top two waveforms show the side and center panel luma lows after raster mapping them into their original widescreen time slots. As before, both signals have one or more "bad pixels" at the extreme edges. Although more "good" pixels apparently exist for splicing, only one "good" overlapped pixel is actually used. The reason is that the side and center panels in each signal have unequal bandwidths; when the luma lows are added to the luma highs, they must be spliced in exactly the same manner to provide a uniform frequency response throughout the entire splicing region.

Figure 17 shows the relevant chroma signals in the decoder. The size of the chroma overlap region allows pixels for splicing. The chroma signals are now abutted. A final lowpass filter for I and Q serves to smooth the seam.

Figure 18 shows the spatial relationship of decoder splices for recovered luma lows, luma highs, and chroma. Note that the luma and chroma splices occur in slightly different locations on the widescreen line.

Scrambling of the Component 2 Signals

In order to minimize visibility of the modulation on the alternate subcarrier, the modulating signals (YSH, QS, IS, and YSD) are scrambled. Scrambling is accomplished by barrel-shifting with pseudo-random cut-points within a scan line. The scrambling sequence is 16 fields long. The barrel-shifting technique is similar to the technique used for pan-and-scan[33].

ACTV's scrambling is described in Figures 19 and 20. The edges of the right and left side panels are tapered (to preserve frequency response and avoid edge effects, as has been described previously) and then spliced together. The pixels in the spliced region are discarded at the receiver. A random cut point is inserted and the pixels are barrel-shifted around it. Line-pairs, that are intraframe-averaged must have the same cut point. Tapers are added to the new edges of the rotated line. This signal is expanded and modulates the alternate subcarrier.

Noise Reduction

The purpose of noise reduction is to reduce the noise power in the recovered side panels to a level that perceptually matches that of the main panel. Since the side panels are transmitted in subchannels that are physically or perceptually hidden from view, special noise reduction processing can be applied to them.

The basic technique for ACTV's noise reduction is the non-linear process of companding, shown in Figure 21. An estimate of each pixel is made at the encoder using a predictor. For YSL, the predictor is motion-adaptive. This estimate is used to select one of a multiplicity of non-linear compression characteristics. These non-linearities have the net effect of attenuating the received noise power when the input signal is substantially similar to its estimate. At the decoder, an estimate of each pixel is made using a predictor. This estimate is used to select an optimal expansion curve for each pixel. The expanded value is used as the recovered, noise-reduced pixel.

Companding can either improve or degrade the S/N of the recovered signal; to be effective, the signal amplitude must, with high probability, be in the noise reduction portion of the curve. ACTV tailors the compandor to the recovered signal by using prediction at both the encoder and decoder to select among different companding curves.

Figure 22 shows how both low-level and high-level signals are noise-reduced by placing maximum compandor gain at the predicted signal level. Encoder and decoder curves are complementary. An example family of companding curves is shown in Figure 23. A schematic of the complete system is shown in Figure 24. The indicated noise is that which is added in the transmission path. Note that the encoder includes a copy of a typical decoder circuit so that distortions introduced, for example, by channel filters, are compensated before transmission.

The slope of the companding curve is described by its "mu-law." Figure 25 shows an expander function which is a mu-law of 30. Typical probability density functions (PDF's) for input and output noise are shown. The sigma of the output noise can be expressed:

$$\text{SIGMA_OUT} = [\Sigma_x \text{ PDF_IN}(x) \text{ EXPANDER}^2(x) \, \Delta x]^{0.5},$$

where SIGMA_OUT is the RMS output noise and Δx is the step size used in the summation. This SIGMA_OUT can be compared with SIGMA_IN to determine the amount of noise reduction provided by the expander.

Figure 26 shows the resultant noise reduction of mu-law companders for subchannel SNR from 0 dB to 50 dB; the subchannel SNR is defined as maximum peak-to-peak signal/RMS noise - standard video SNR definition. It must be stressed that the subchannel SNR is the SNR with which the noise reduction circuit is operating-- it may be lower than the main channel ("NTSC") SNR - particulary when signals are attenuated. Figure 26 shows that the amount of noise improvement, in general, goes down at low SNR. This is because the input signal falls outside of the maximum noise gain region more and more with increasing noise. At extremely low SNR, the amount of noise improvement rises again. This is because the noise gets clipped by the expander function (notice in Figure 25 how inputs having absolute value greater than 0.5 are clipped at the output to -0.5 or +0.5).

The ACTV subchannels operate at an input S/N that requires an approximately 10 dB improvement from noise reduction.

Figure 26 allows predictions about noise reduction requirements. In the absence of noise reduction, the difference between the channel SNR and any subchannel SNR is a fixed amount [dB]. Noise densities have been calculated for the various subchannels, in order to determine noise reduction requirements. The relation between the channel and subchannel SNR is determined by integrating their noise power density functions. The relevant noise density functions are those that appear at the input to the noise reduction decoder circuits.

Progressive Scan Conversion

The chroma signals are converted from interlace to progressive scan format using non-adaptive vertical line averaging. The luma signal, however, is converted using a motion-adaptive algorithm. For each interpolated pixel, a motion detector produces a k-signal, which is used to soft-switch between a temporal and vertical estimate. The k-signal is derived by first generating seven absolute differences above, below, and before the interpolated pixel. Four of these differences are vertical-temporal, three are pure temporal. The maximum of these absolute differences is computed, horizontally lowpass filtered, and spread both spatially and temporally. Finally, the k-signal is normalized between 0 (no motion) and 1 (fast motion) and is used to switch smoothly between a field-repeat (no motion) and vertical average (fast motion).

Pan-and-Scan

In all the foregoing, we have assumed that the 4x3 panning window is horizontally centered on the widescreen image. However, when the main action is off-center, it may be desirable for the compatible 4x3 NTSC display to follow the action by moving the panning window. This is known as pan-and-scan.

Smooth pan-and-scan capability is implemented in ACTV by redefining what the encoder and decoder process as "center" and "side" panels. This is done by performing a *barrel-shift* on the widescreen input signals followed by an *inverse barrel-shift* on the decoded widescreen output signals. A panning offset is generated for each frame of the widescreen source material. The panning offset is a single number that indicates the horizontal offset of the panning window, in widescreen pixels. When the panning window is centered, the offset is zero. When the panning window is left of center, the offset is negative, and so on. The panning offset is transmitted to the decoder as a digitally encoded signal

on one line in the vertical blanking interval. To work properly with intraframe averaging, the panning offset can change only on frame boundaries.

At the encoder, the panning offset is used to barrel-shift the YIQ widescreen source image in such a manner as to keep the image in the panning window centered on the active line. The side panels remain hidden even when the 4x3 portion of the picture is panned left or right. As shown in Figure 27, when the panning window is shifted left, the widescreen source is shifted right; the portion of the right side panel that falls outside the active line is cut and pasted into the gap next to the shifted left side panel. The barrel shift can be viewed as a preprocessing operation that centers the 4x3 panning window on the active line. The rest of the encoder operates independently of pan-and-scan; it assumes a centered "center panel" and equal-width "side panels". The barrel-shift method of pan-and-scan redefines "center panel" to mean the image in the panning window. At the decoder, the widescreen components are recovered, and an inverse barrel shift is performed to realign the image spatially.

ACTV's pan and scan capability will not be implemented in the hardware delivered for testing.

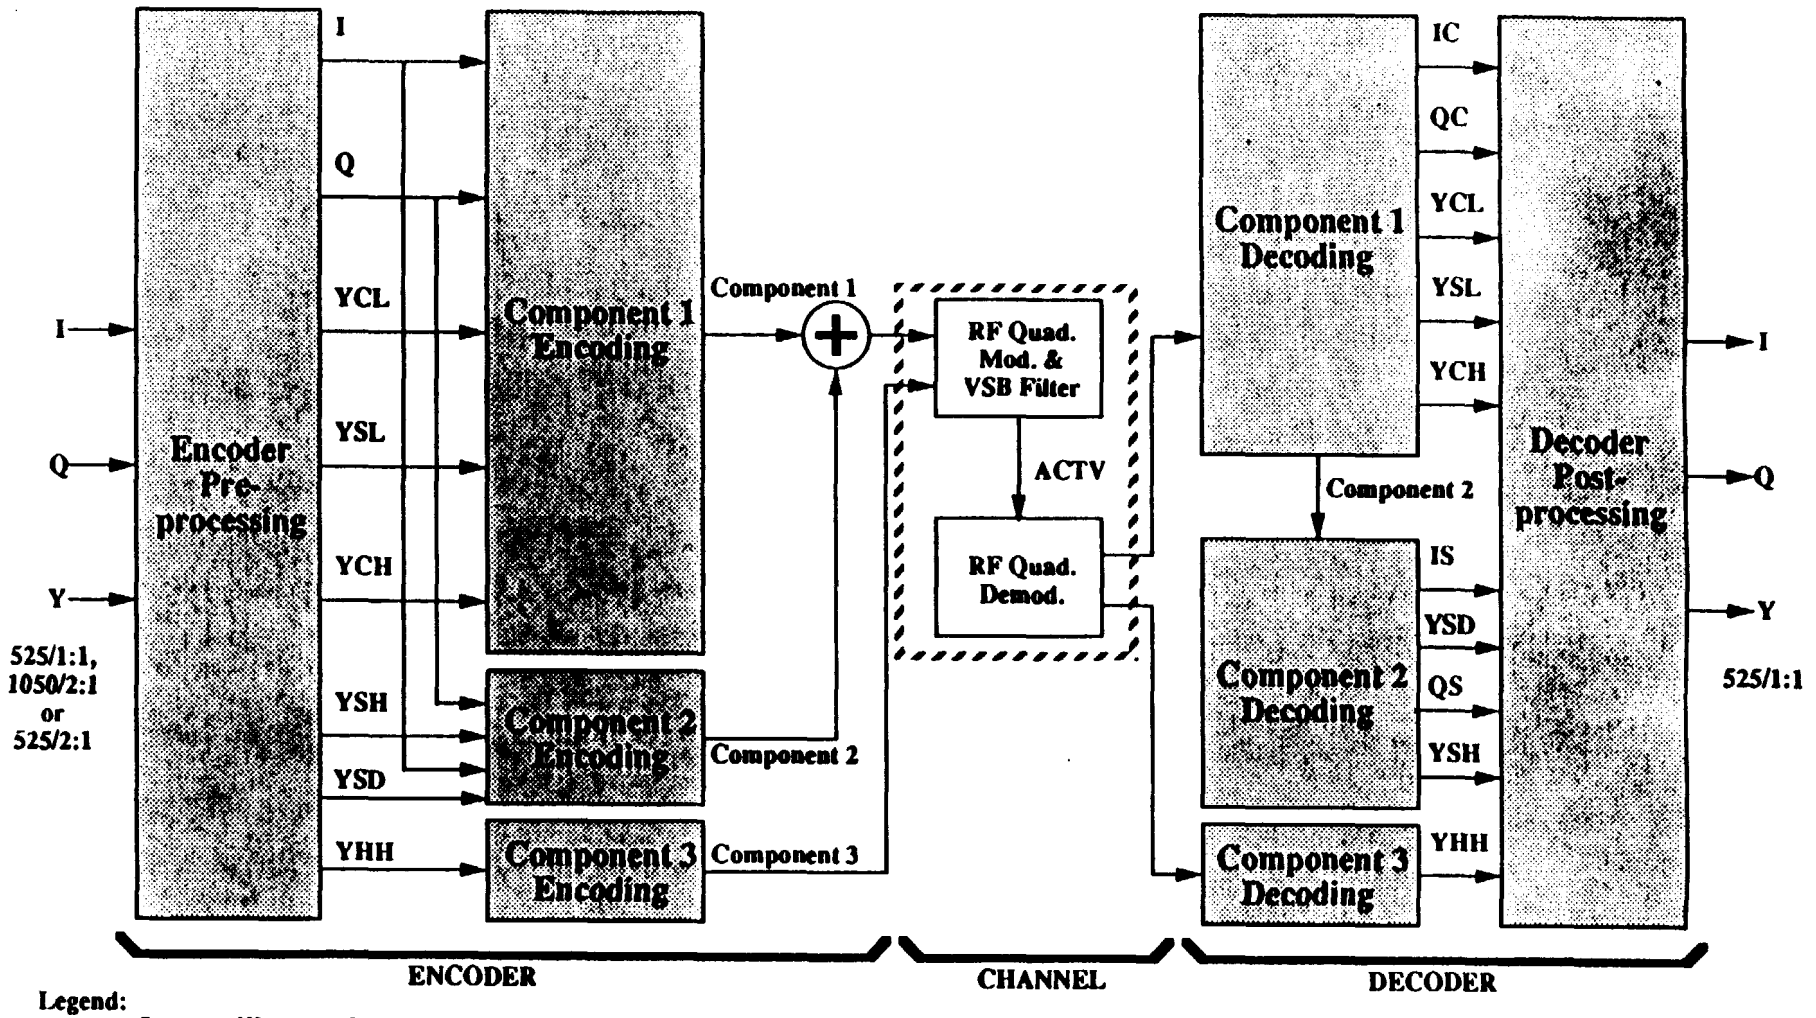

Legend:

- I: I color-difference signal
- IC: I center panel
- IS: I side panel
- Q: Q color-difference signal
- QC: Q center panel
- QS: Q side panel
- Y: Y luminance signal
- YCL: Y center panel lows
- YCH: Y center panel highs
- YSL: Y side panel lows
- YSH: Y side panel highs
- YHH: Y high-highs (extra detail)
- YSD: Y side difference

Fig. 1 ACTV System

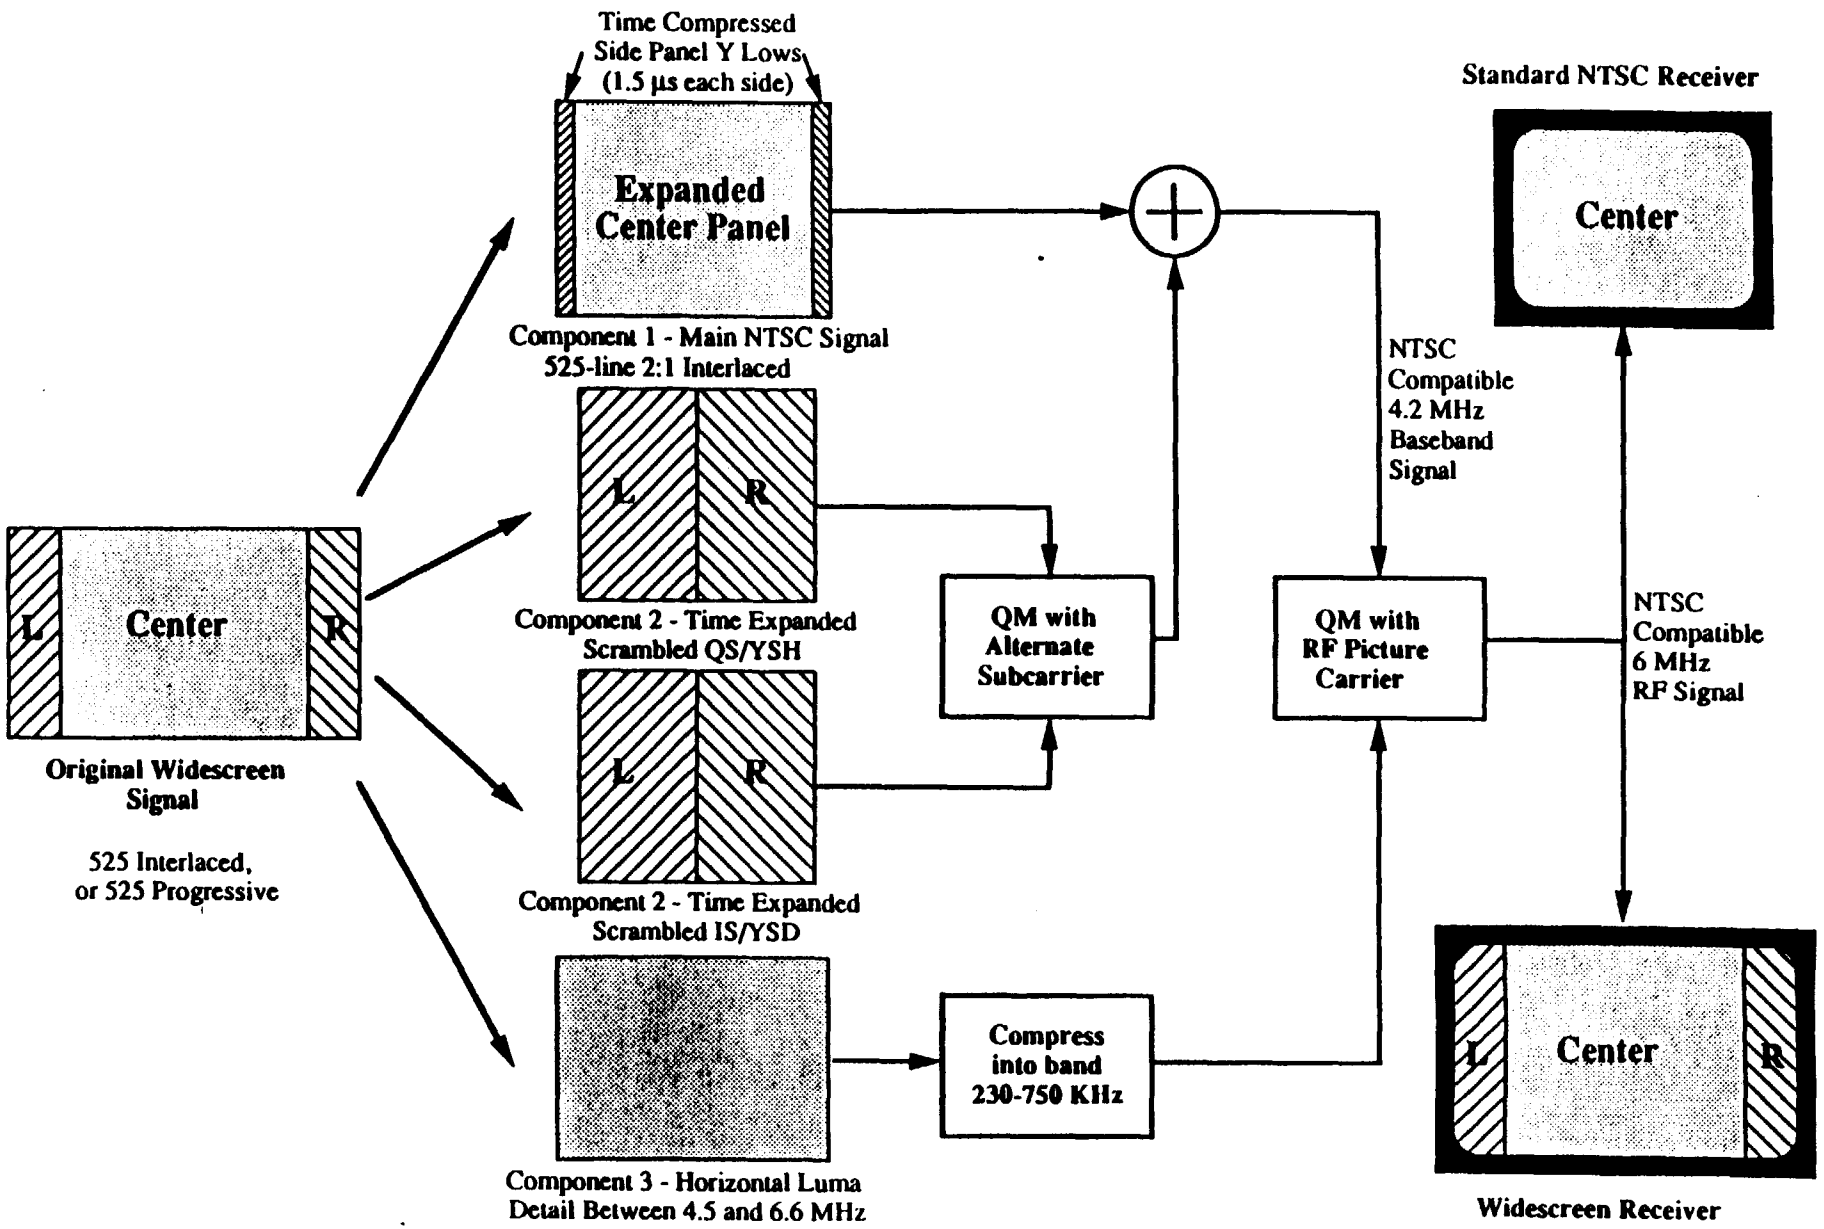

Fig. 2 ACTV encoder

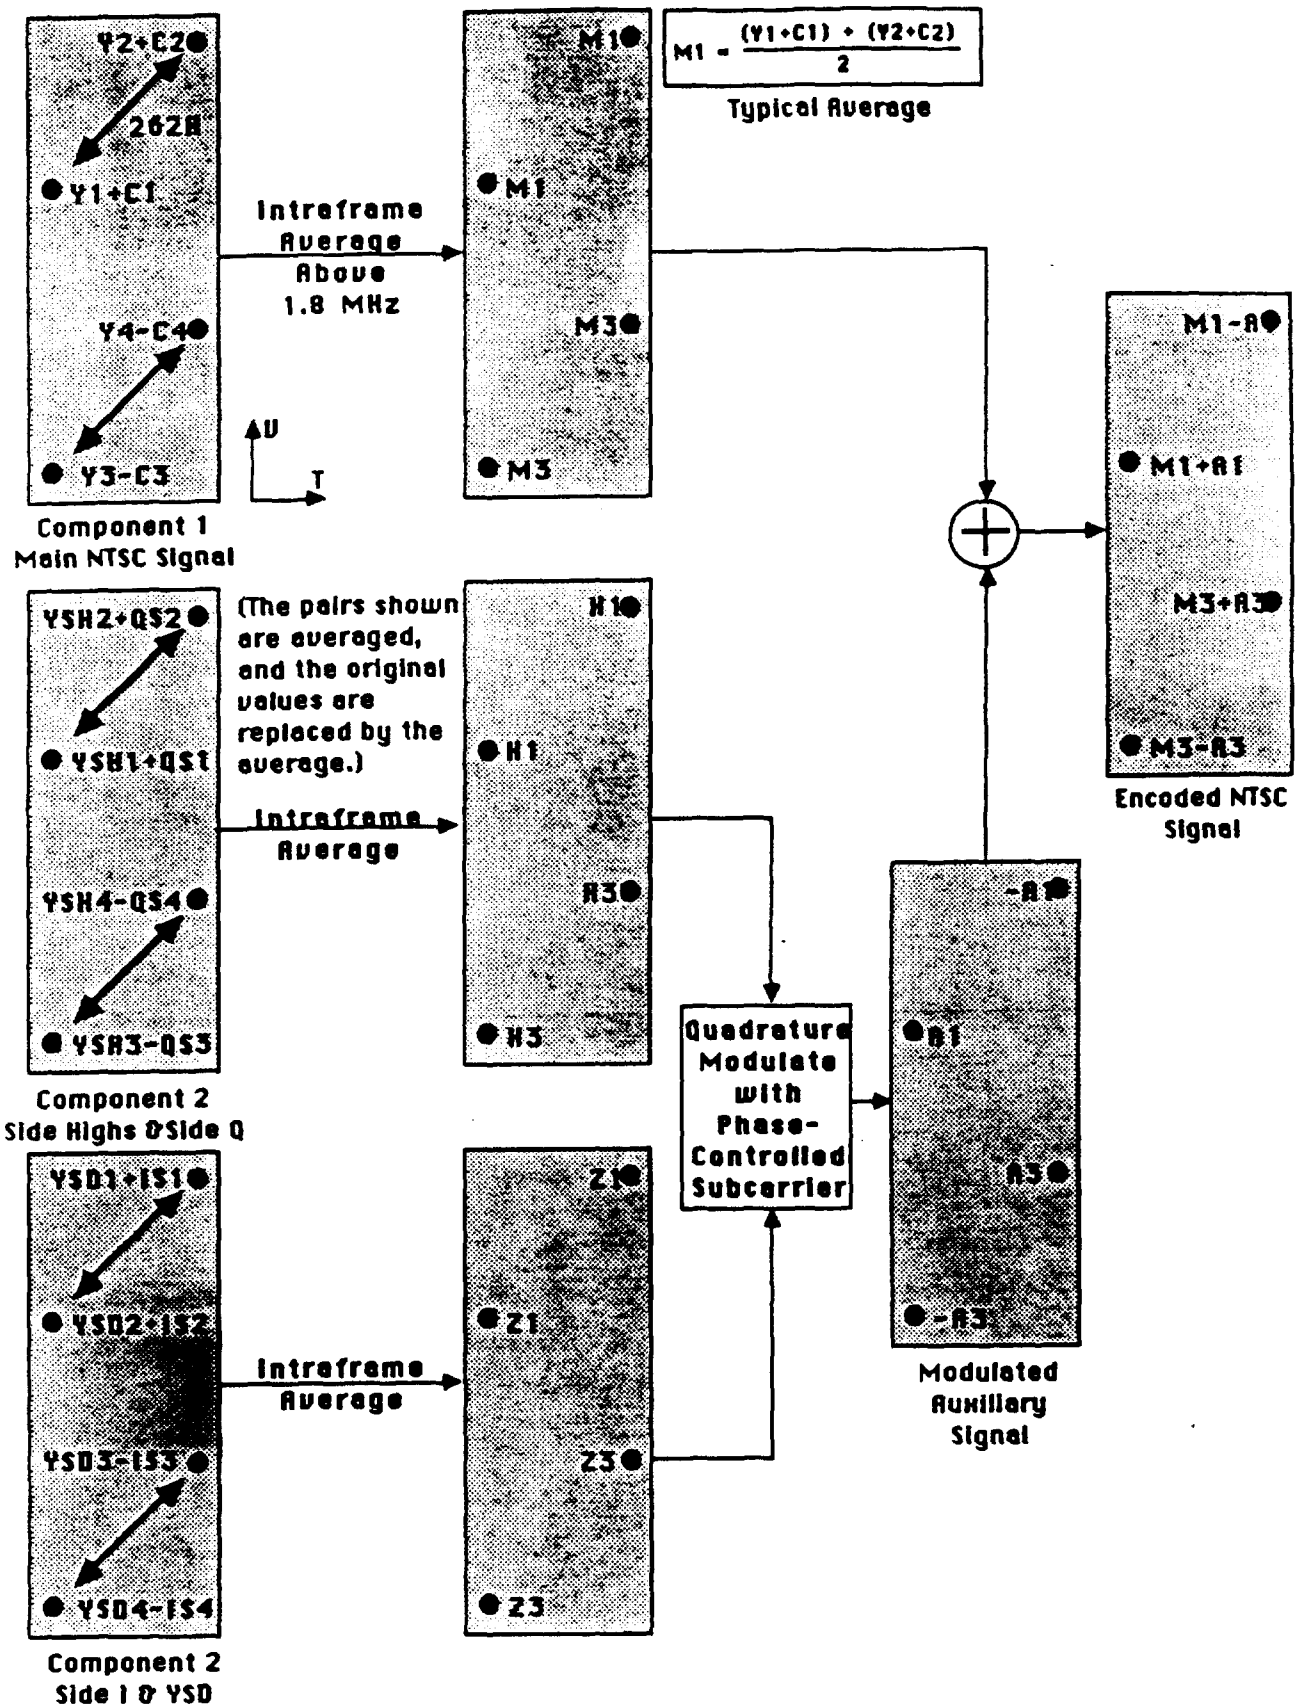

Fig. 3 Graphical explanation of intraframe averaging.

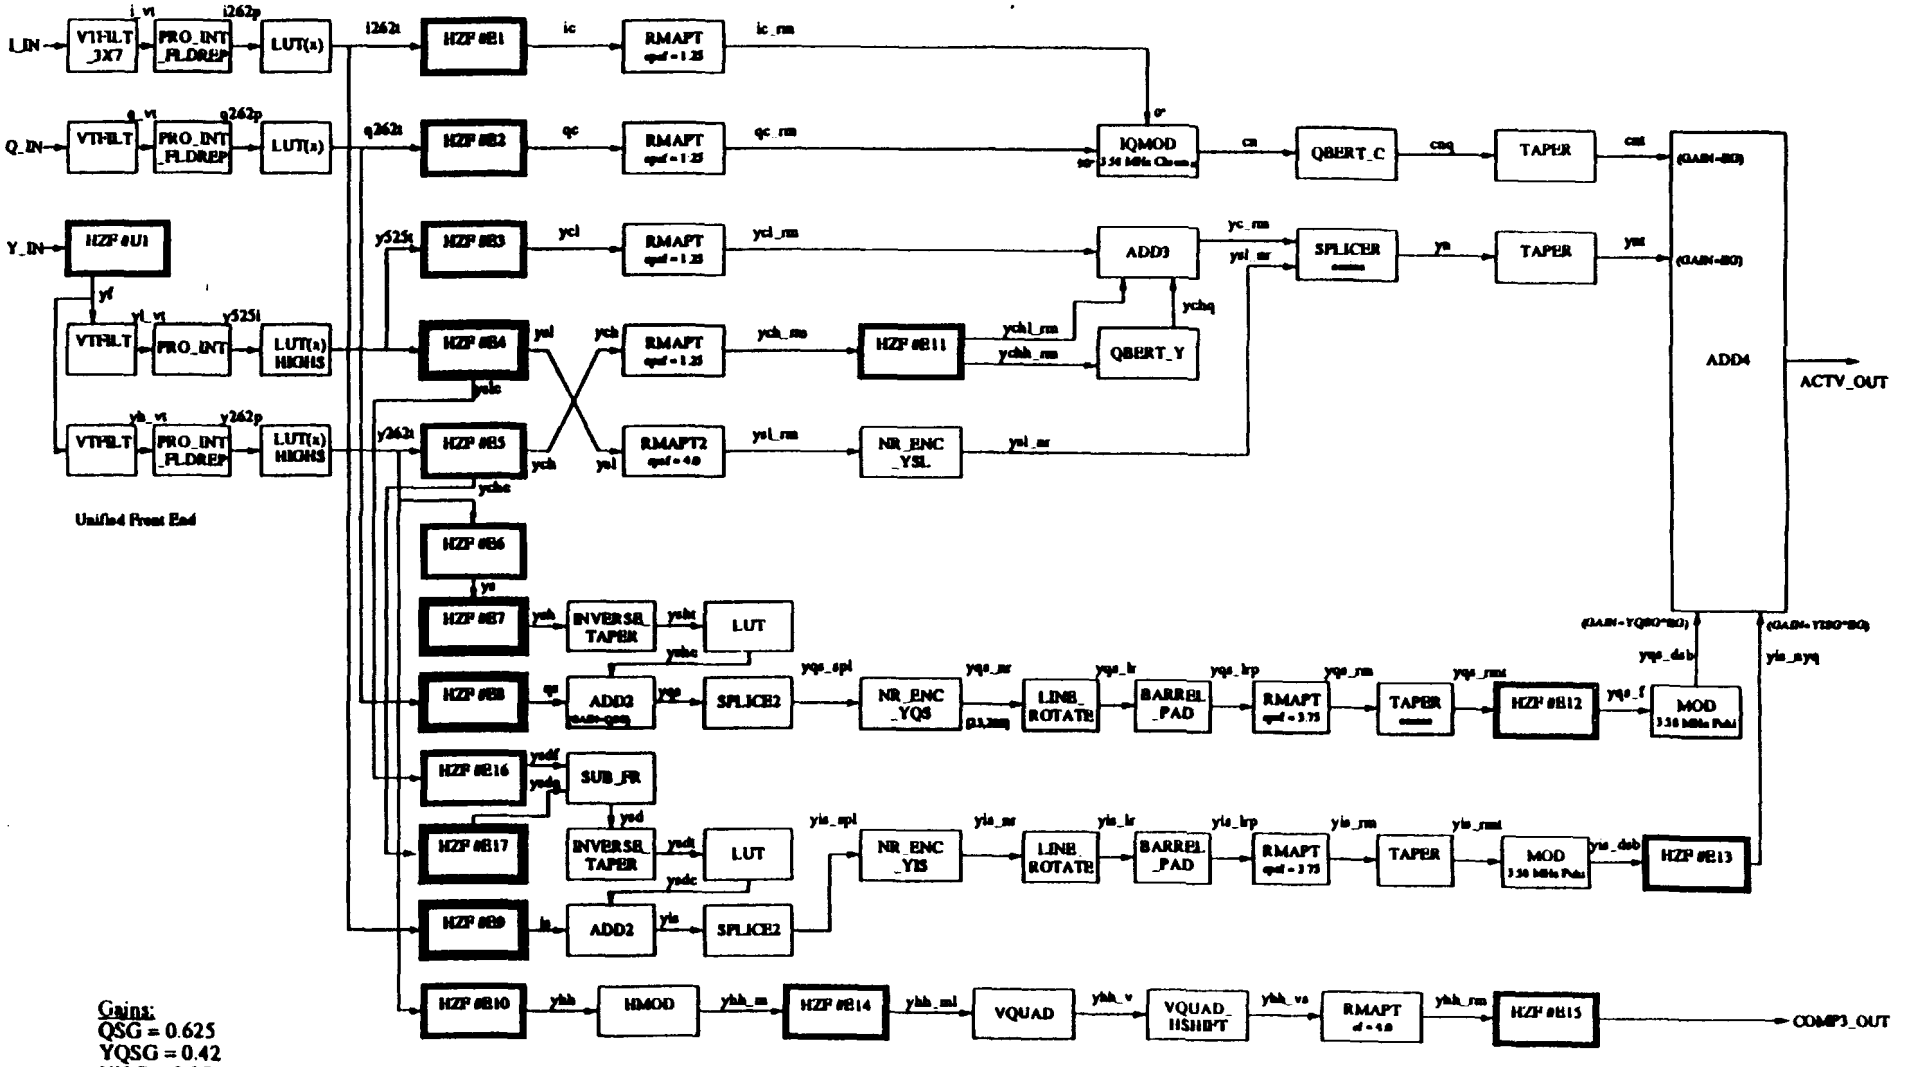

Gains:
QSG = 0.625
YQSG = 0.42
YISG = 0.35
EG = 0.9

LUT(x) is function of horizontal position

Fig. 4 Encoder

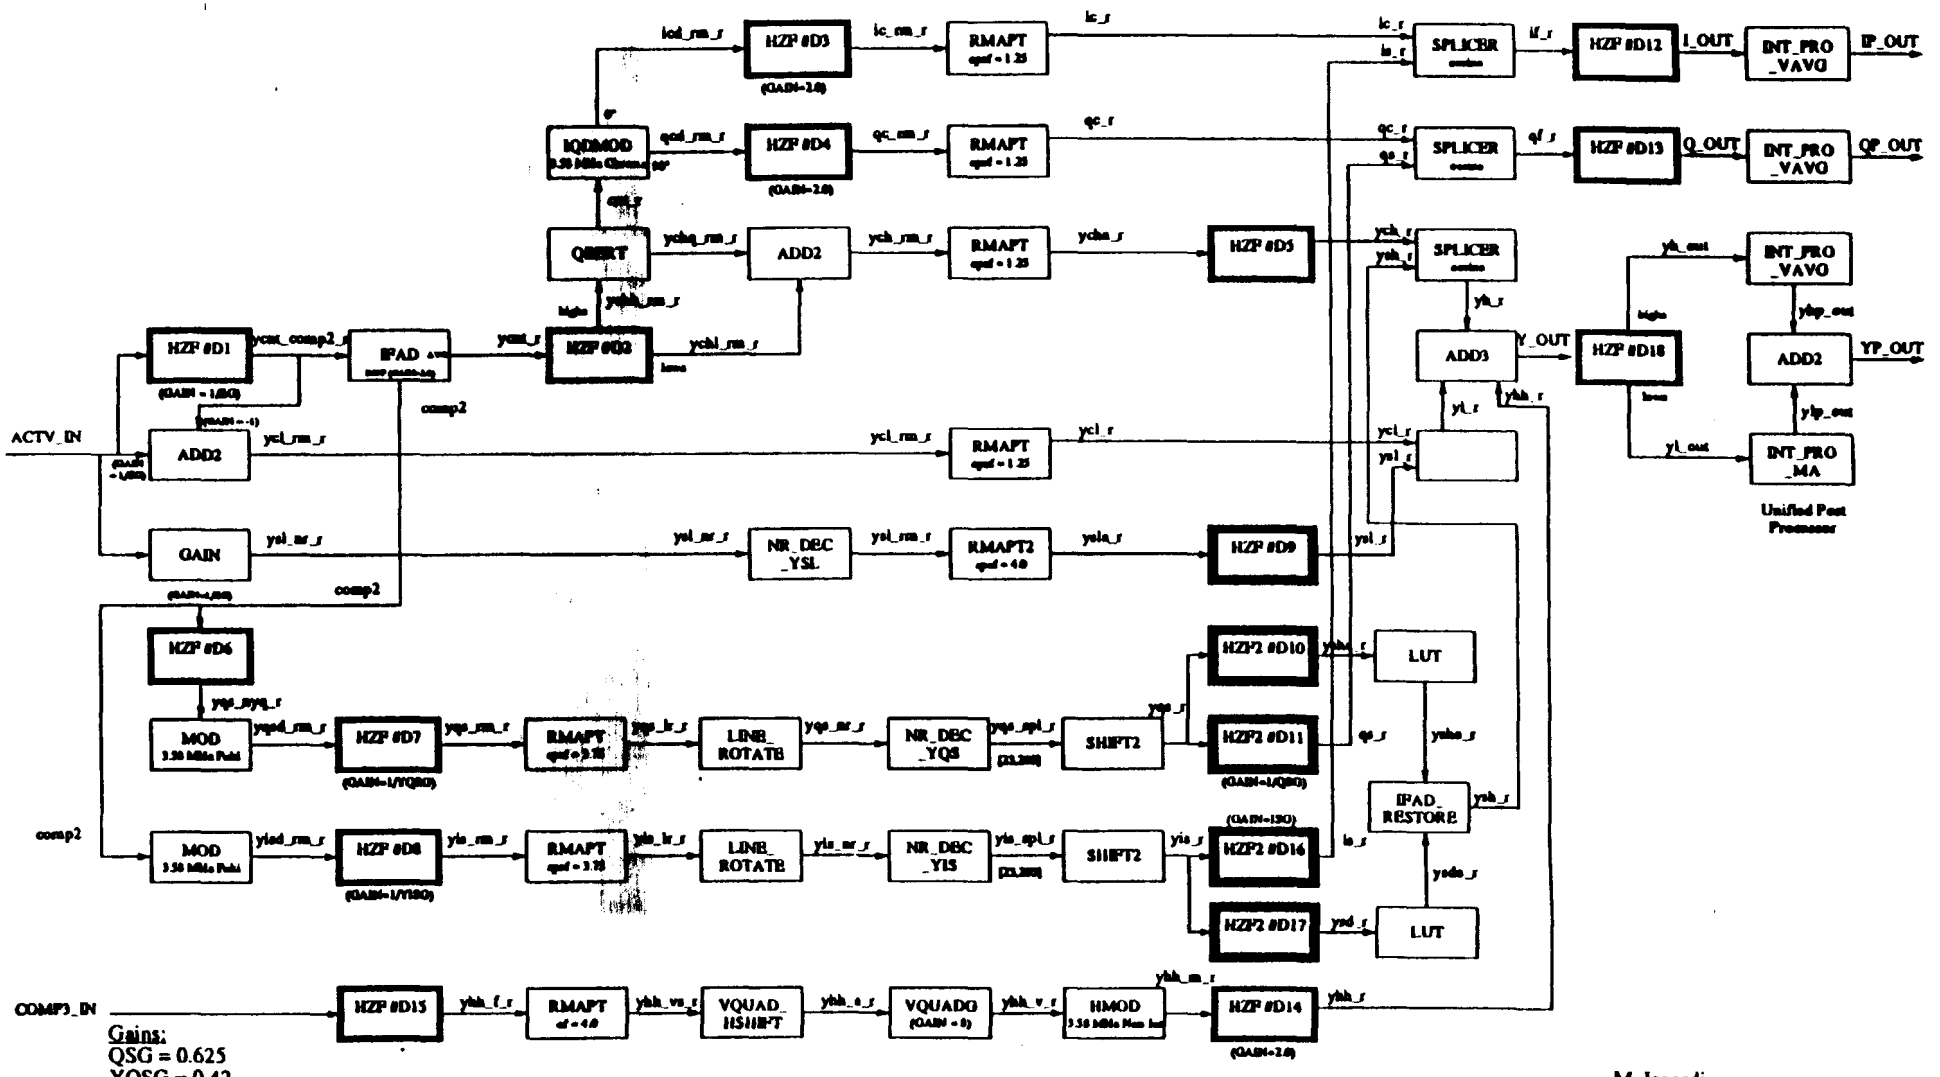

Fig. 5 Decoder

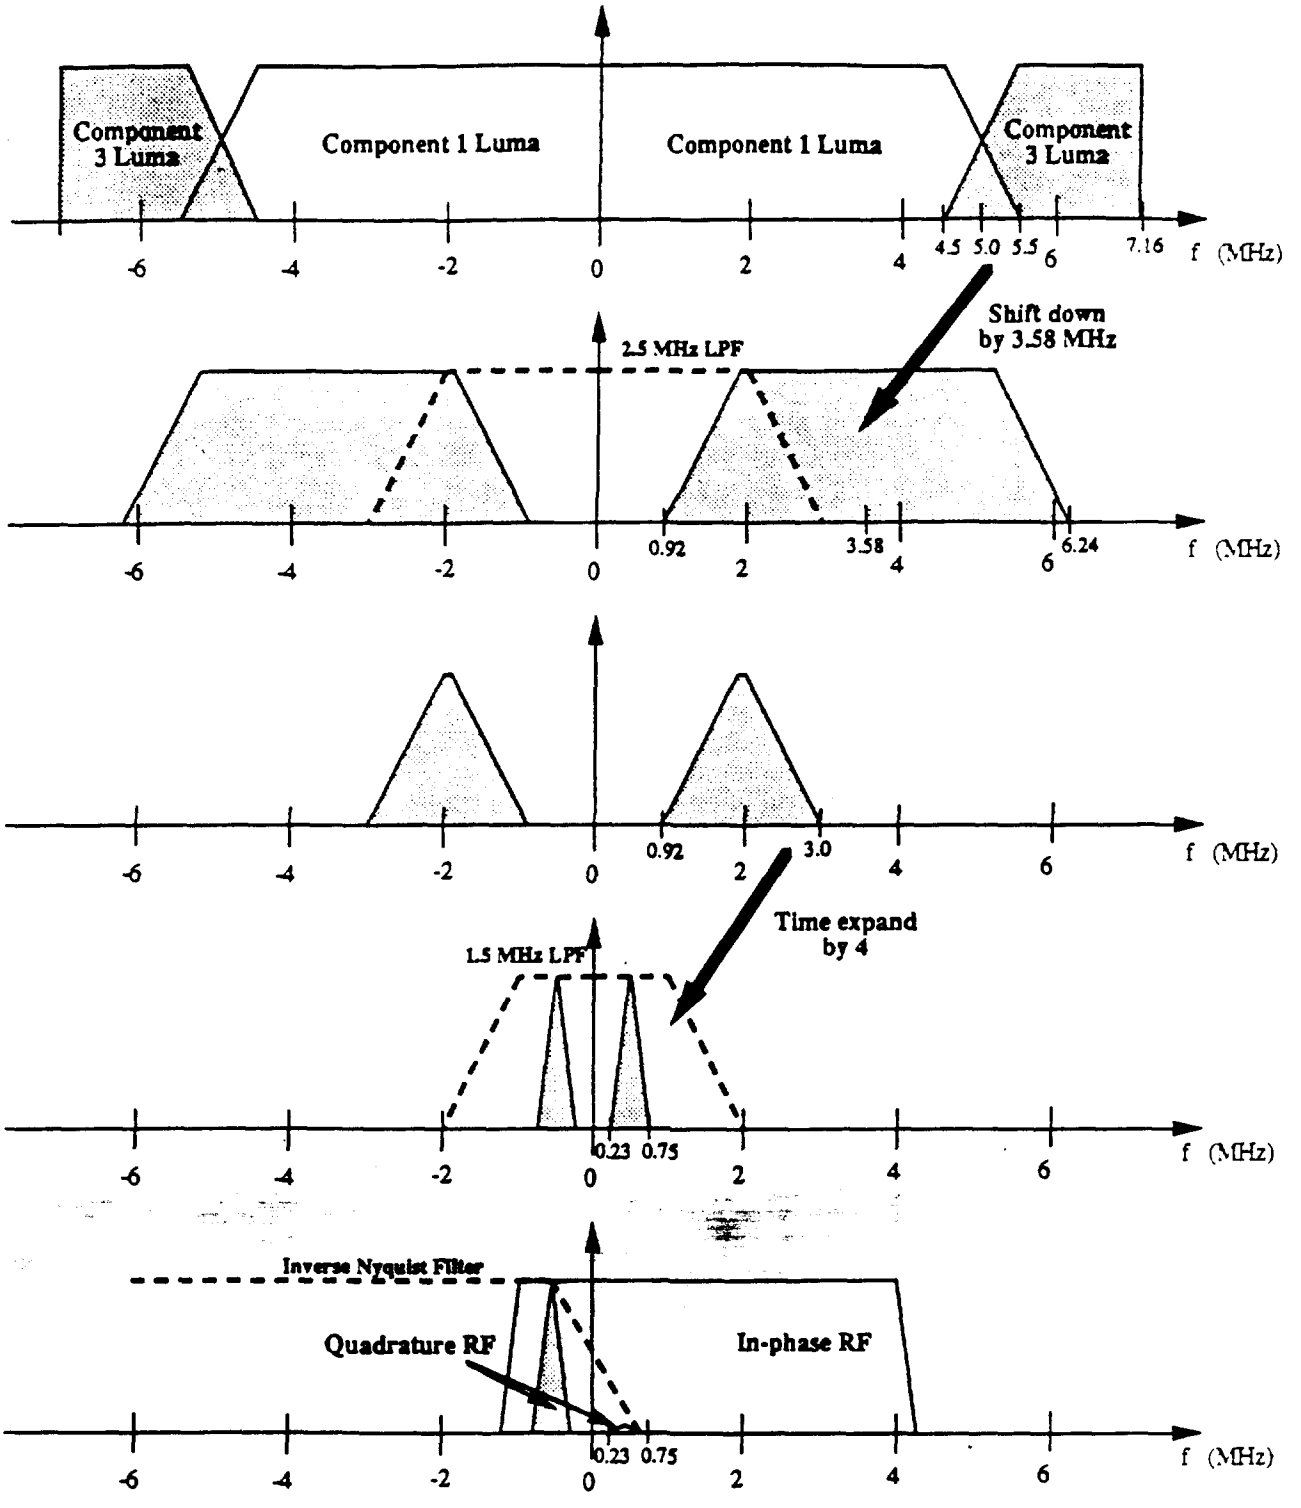

Fig. 6 Encoding of Component 3 in RF Quadrature

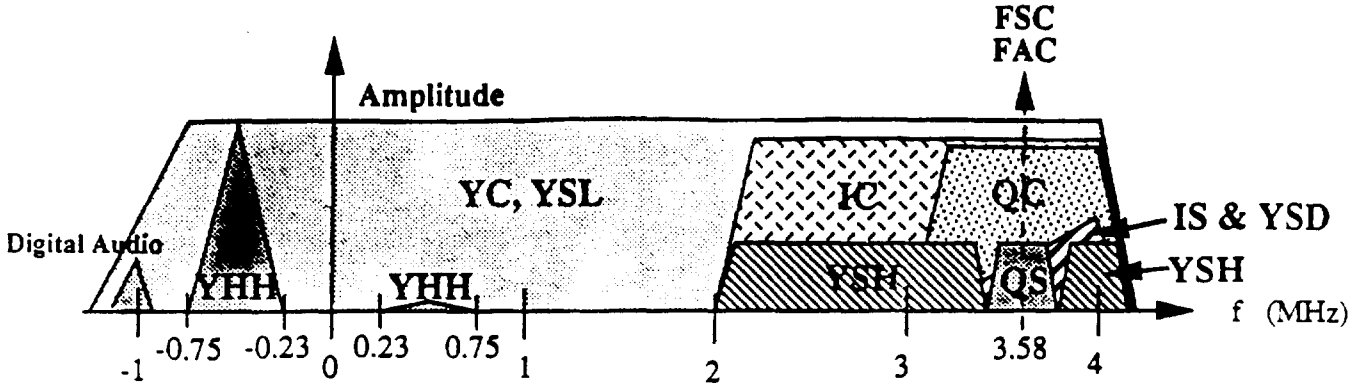

(a) RF spectrum. Amplitudes not to scale.

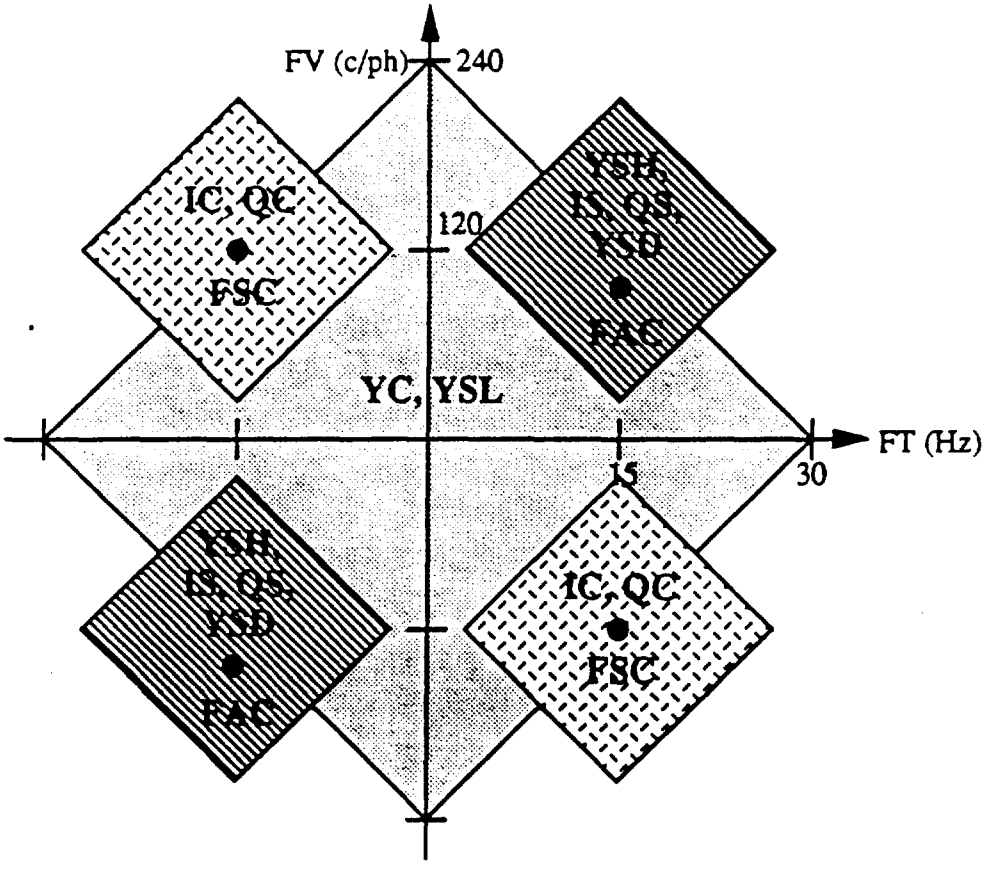

(b) V-T spectrum for horizontal frequencies in the range 2.0 - 4.2 MHz.

Fig. 7 RF and V-T Spectra of ACTV Signal

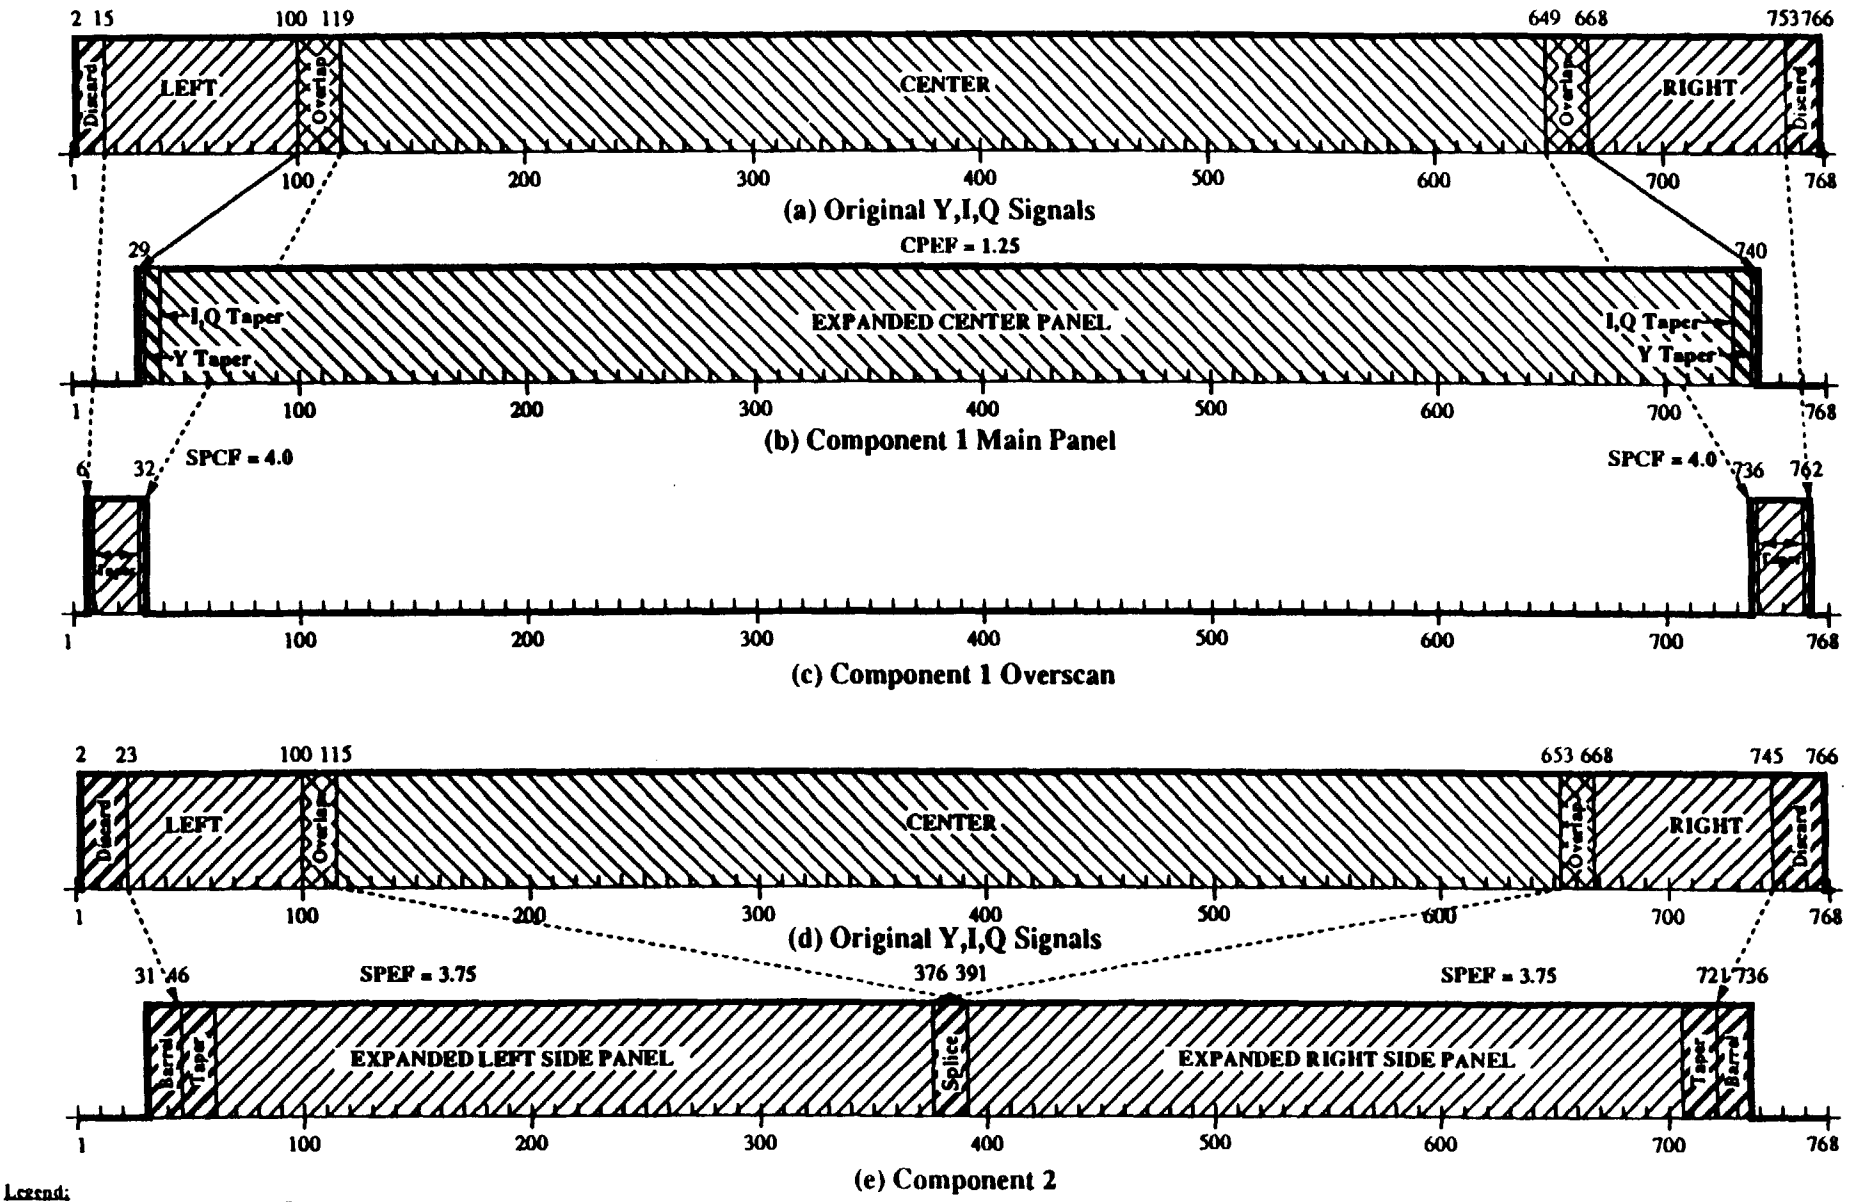

Legend:
CPEF: Center Panel Expansion Factor
SPCF: Side Panel Compression Factor
SPEF: Side Panel Expansion Factor

Fig. 8 Raster Mapping Operations - Encoder

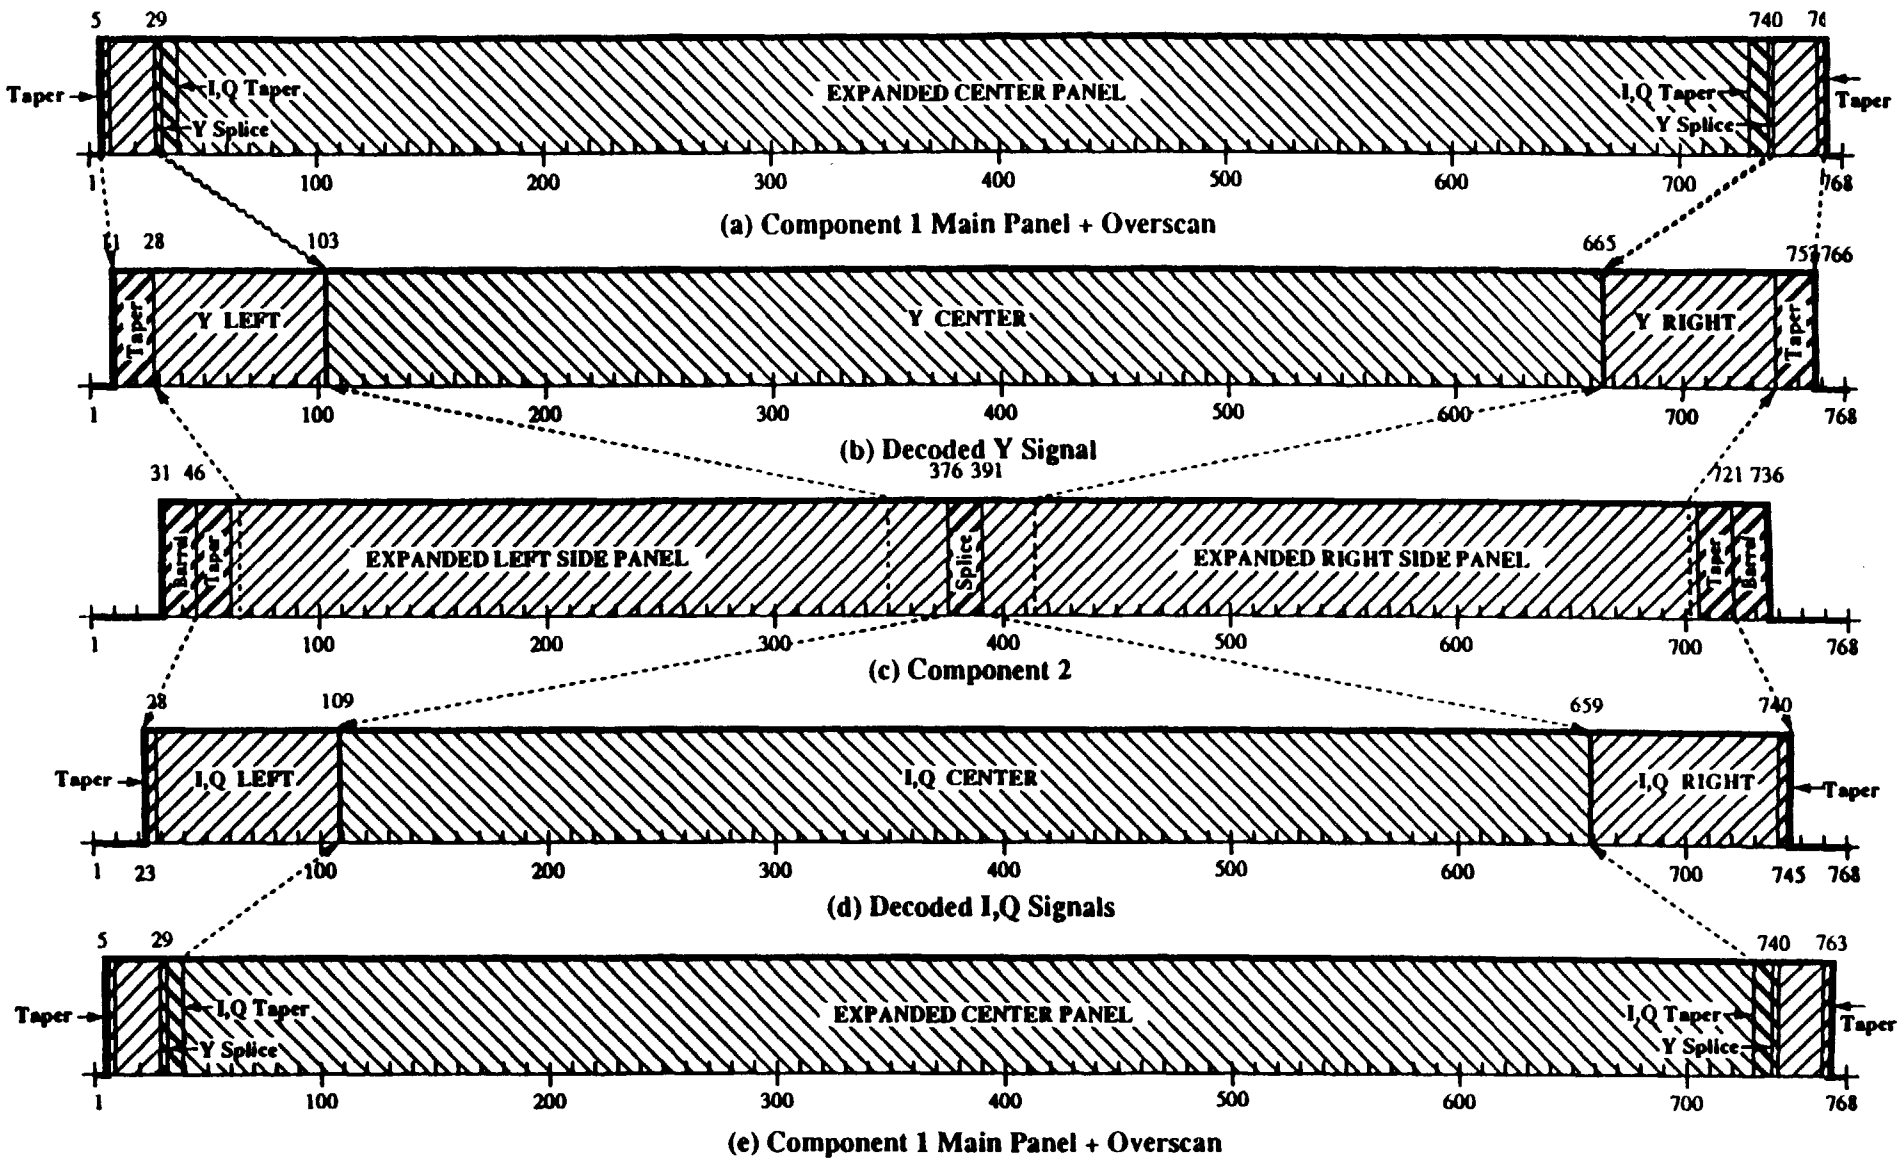

Fig. 9 Raster Mapping Operations - Decoder

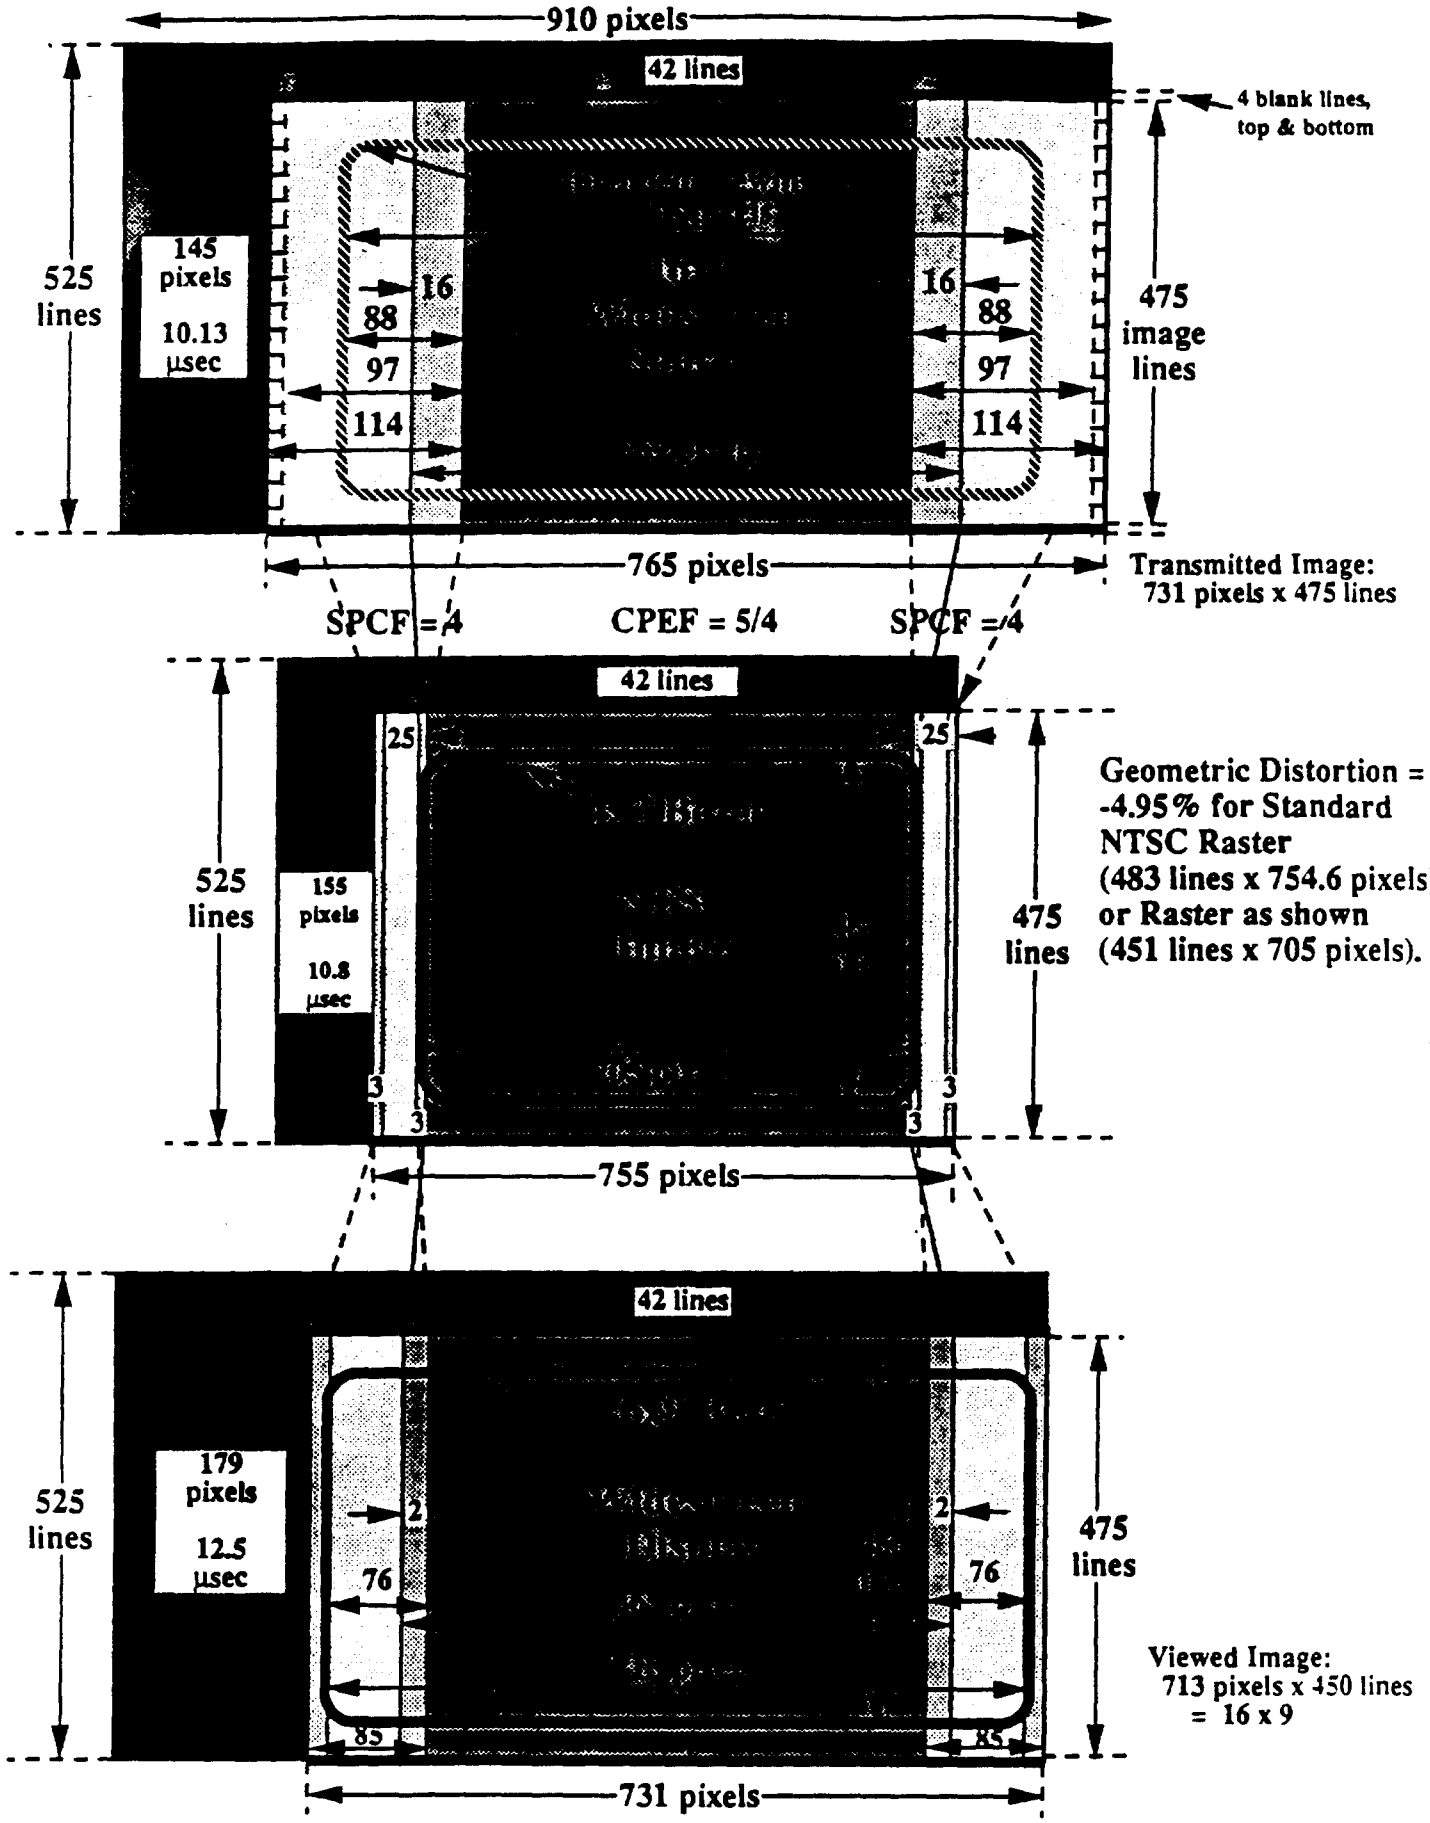

Fig. 10 Raster Parameters for Component 1

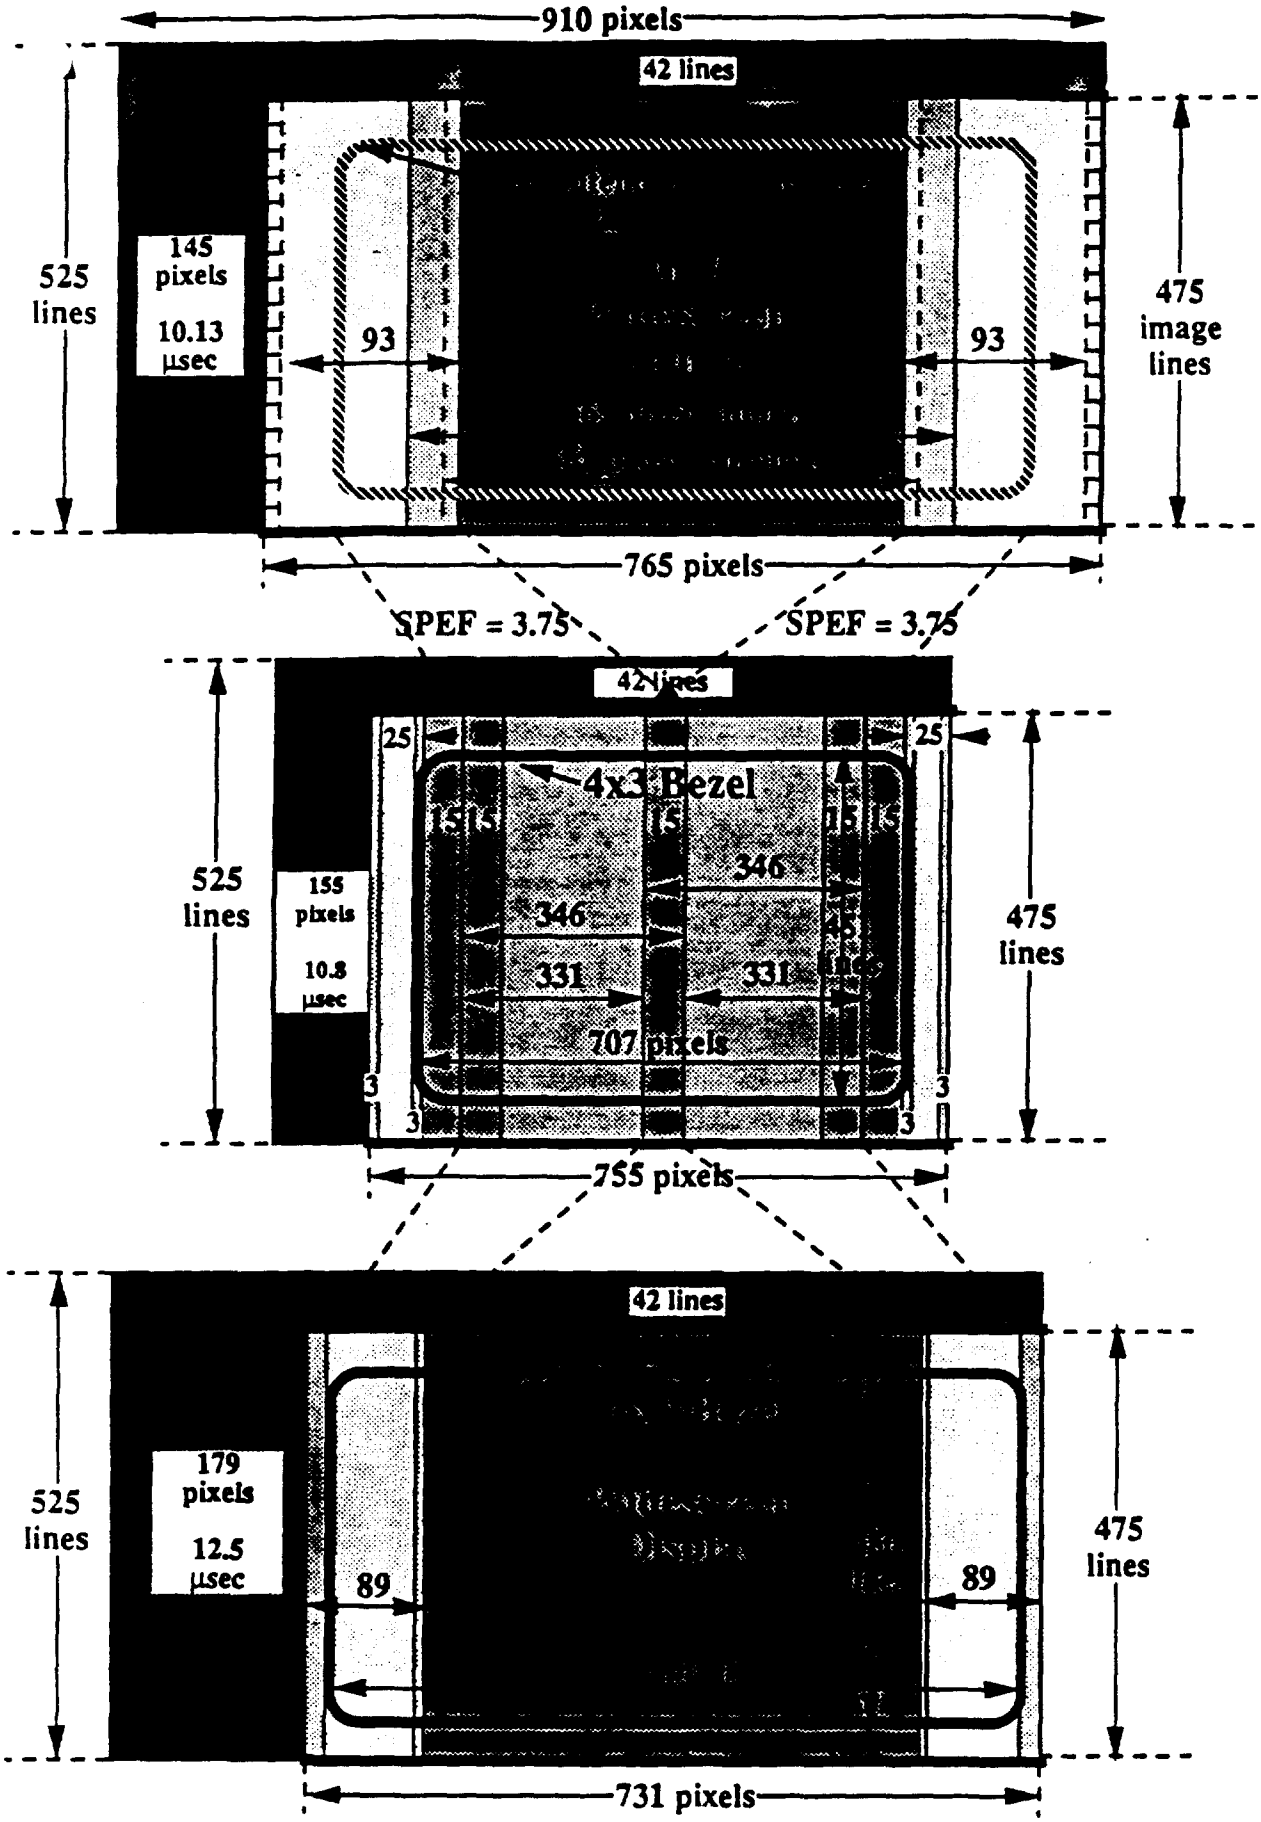

Fig. 11 Raster Parameters for Component 2

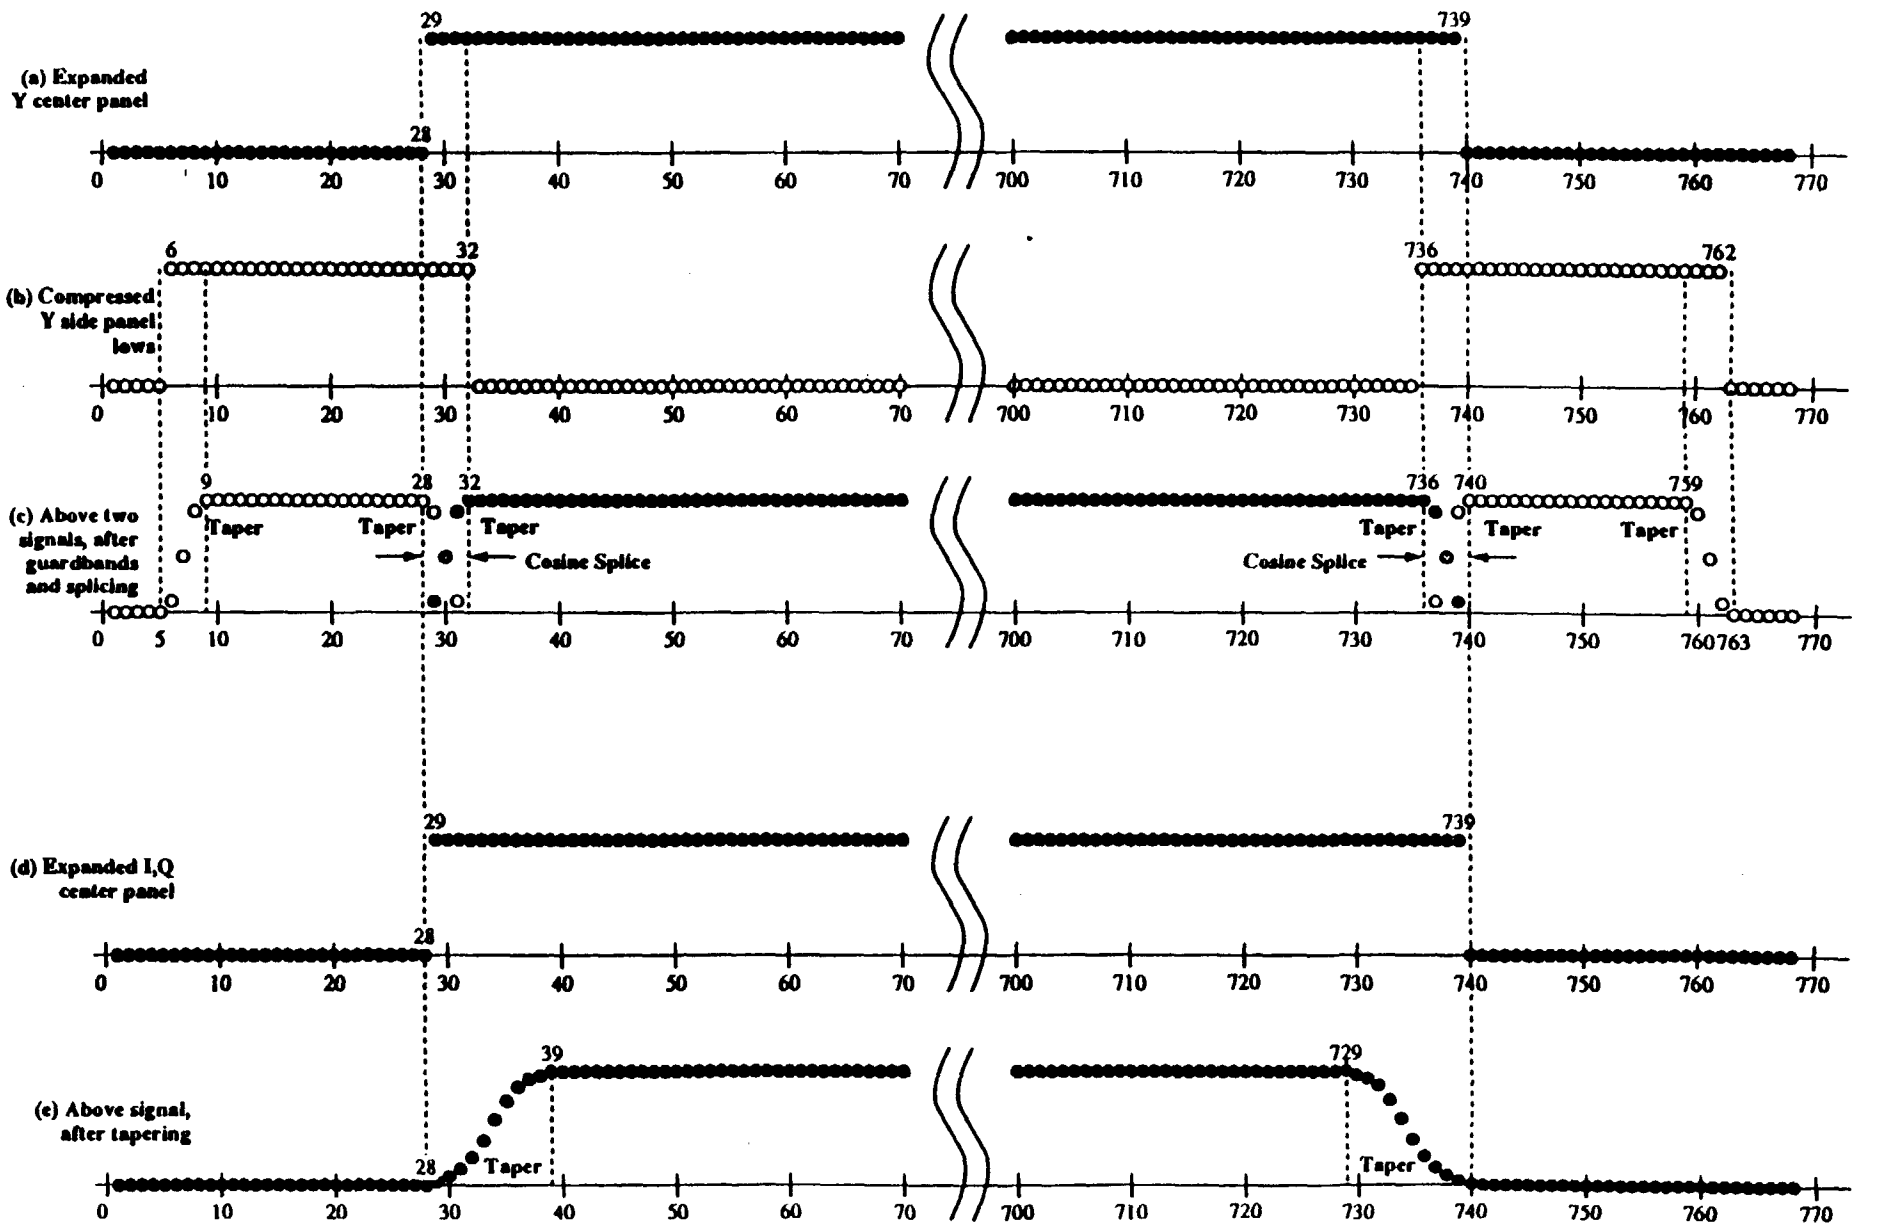

Fig. 12 Component 1 Guardbands

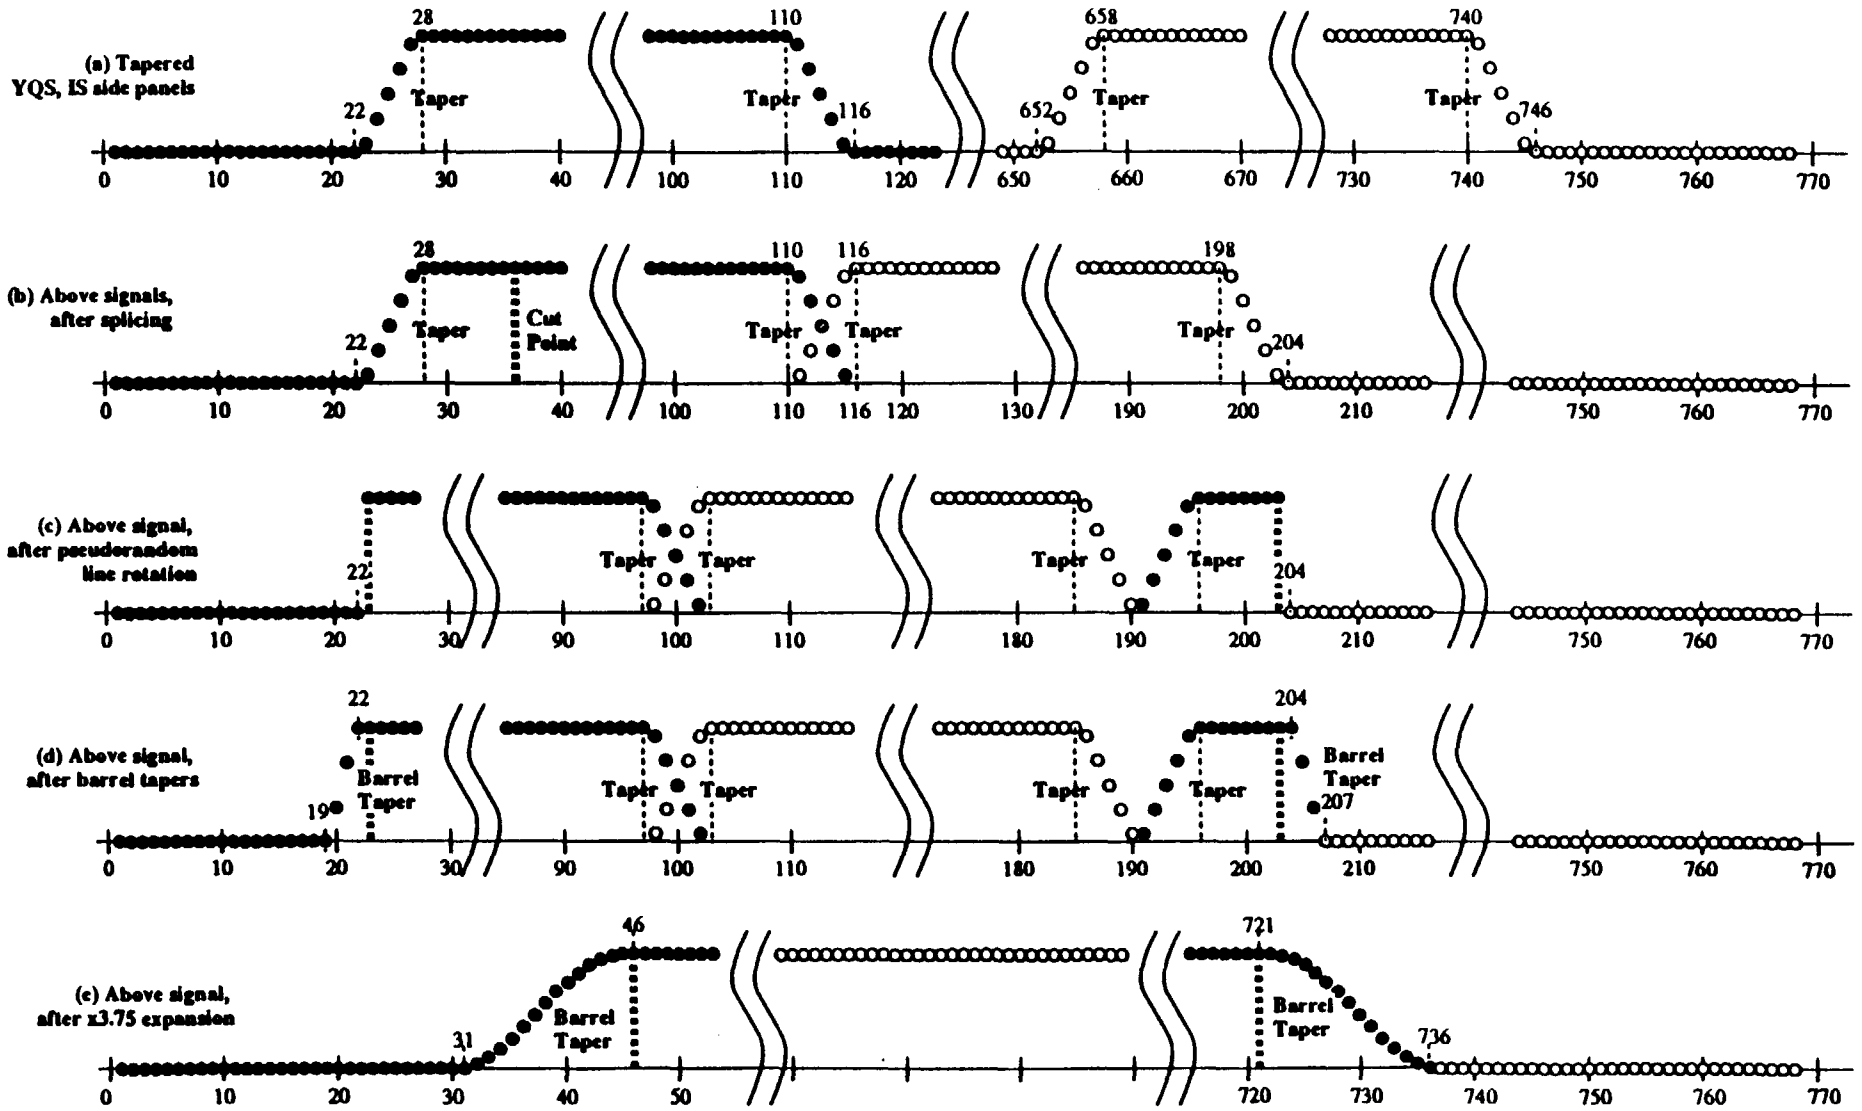

Fig. 13 Component 2 Guardbands