COLOR DATA INTERCHANGE: TECHNOLOGY AND STANDARDS

Robert R. Buckley*

Abstract: Displaying color images in different media implies the interchange or transmission of color image data. Because transmission utilizes a limited resource (channel bandwidth or capacity), a great deal of effort has been invested in the efficient representation or coding of color images for transmission and storage. The current device-independent or CIE-based standards for color interchange will be reviewed, along with the algorithms for coding image data. It will be noted that current approaches are traceable to a few key ideas that appeared in connection with the emergence of color television around forty years ago.

INTRODUCTION

This presentation will review the technology and standards for color data interchange. Because images stress systems more than other color data, it will focus on color image interchange. Besides reviewing technology and standards, I will also describe some history, as many of the approaches used today had their origins in the late 1940's and the 1950's as a result of the new opportunities and possibilities created by color television.

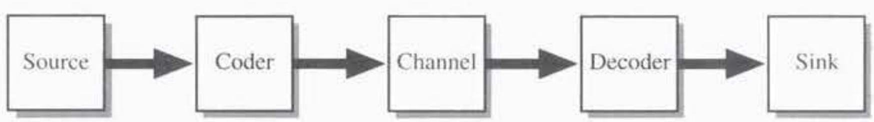

Figure1. Image Communications Model

In its simplest and most general form, a system for communicating or interchanging color image data consists of a source, a channel and a sink or destination (Figure 1). In color television broadcasting for example, the source is a camera, the channel is a frequency spectrum band, and the sink is a human observer viewing a home receiver. In a color printing system, the source can be a scanner, the channel could be a local area network or a diskette, and the destination could be a color printer.

* Xerox Webster Research Center, Webster, NY

In Figure 1, the coder is the interface between the source and the channel. It transforms the source data into a representation that is suitable for the channel, which is the shared, limited resource. Its use is subject to mutual agreements or standards, designed to enable connectivity and increase capacity. The coded representation is based on constraints imposed by the channel and on how the data will be used at the sink. As a rule, the decoder is the inverse of the coder and faithfully recreates the source data at the destination.

I am going to review three systems for image coding: color television, a digital color system, and the human visual system, describing technology and standards for the first two and commenting on their relationship to the third. I will conclude with a look at the state of standards for intermedia interchange, where the the source and sink use different media or viewing conditions, so that exactly recreating the source data at the destination is not appropriate.

Color Television

In an analog color television system, the source is an RGB color camera, the channel is a 6 MHz band of the frequency spectrum (the amount originally allocated for black-and-white television transmission), and the sink is an RGB display viewed by a human observer. The obvious solution is to transmit RGB signals, and for a time in the early 1950's, the US broadcast standard was the CBS Field Sequential System, based on transmitting and displaying R, G and B fields in rapid succession (McIlwain and Dean, 1956).

If the R, G and B fields each had the same bandwidth as the monochrome field, then the resulting color picture would have the same sharpness and detail as a monochrome picture, but three times the bandwidth. To fit the color signals within the 6 MHz band allocated for transmission, the number of fields per second, the number of scan lines per field, and the video bandwidth along a scan line all had to be reduced. The resulting color picture was noticeably less sharp and had less detail than a monochrome picture.

An alternative proposal for color television broadcast was the RCA Dot Sequential System, in which low-pass filtered R, G and B samples were transmitted in succession along a scan line (McIlwain and Dean, 1956). To increase picture detail, the mixed highs scheme was proposed (Bedford, 1950). It is based on the observation that detail or high frequency information in a picture is practically colorless or achromatic. Therefore, it is not necessary to add three high pass signals–one each for R, G and B–to increase picture detail; adding a single high-passed combination of the three or a mixed highs signal will give the same visual result. Figure 2 shows the block diagram of a mixed highs coder.*

---

* To illustrate the performance of this and other coders, the conference presentation showed several images which are not reproduced here because of space and time limitations.

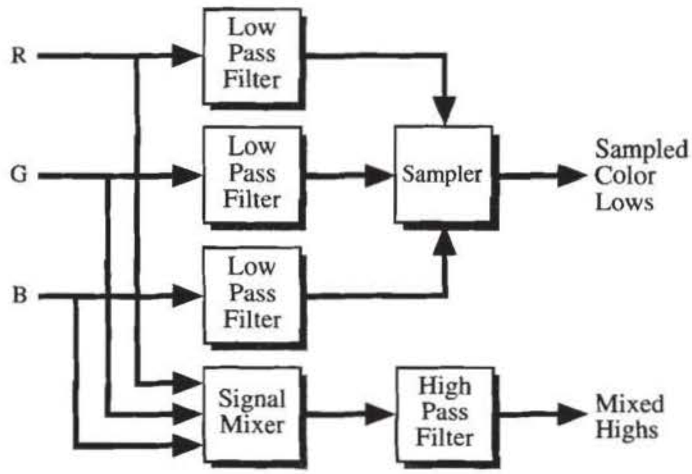

Figure 2. Mixed Highs Coder

In describing the mixed highs coder, A. V. Bedford enunciated the fundamental principle of bandwidth reduction: "avoiding the transmission of information which the eye cannot use." It is as applicable now to the design of digital color systems as it was then to the design of analog systems.

In the mixed highs coder, a high-pass achromatic signal is added to three low-pass color signals. An equivalent arrangement would add two low-pass color signals to an all-pass (low-pass and high-pass) achromatic or luminance signal; this arrangement was proposed by Hazeltine (Loughlin, 1951). Figure 3 shows a block diagram of the Hazeltine constant luminance coder.

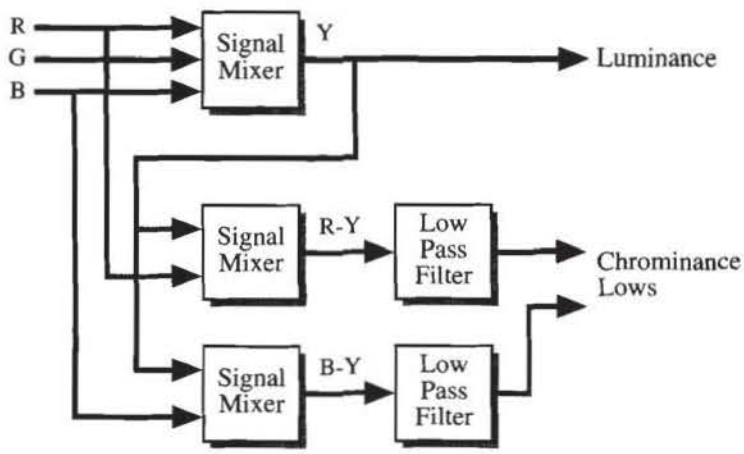

Figure 3. Constant Luminance Coder

The full bandwidth achromatic signal Y is luminance, a linear combination of R, G and B. The signals R-Y and B-Y are called color-difference or chrominance signals; two stimuli with the same luminance component but different chrominance components would have the same brightness but differ in color. Because the spatial frequency response of the eye is narrower for chrominance than it is for luminance, the chrominance signals can be low-pass filtered to save bandwidth.

This coder is called "constant luminance" because the luminance channel is unaffected by the coding of the color information. It was the basis for the NTSC (National Television System Committee) system, adopted in 1953 as the US color television broadcast standard (IRE, 1954). Figure 4 shows a simple block diagram of the NTSC system.

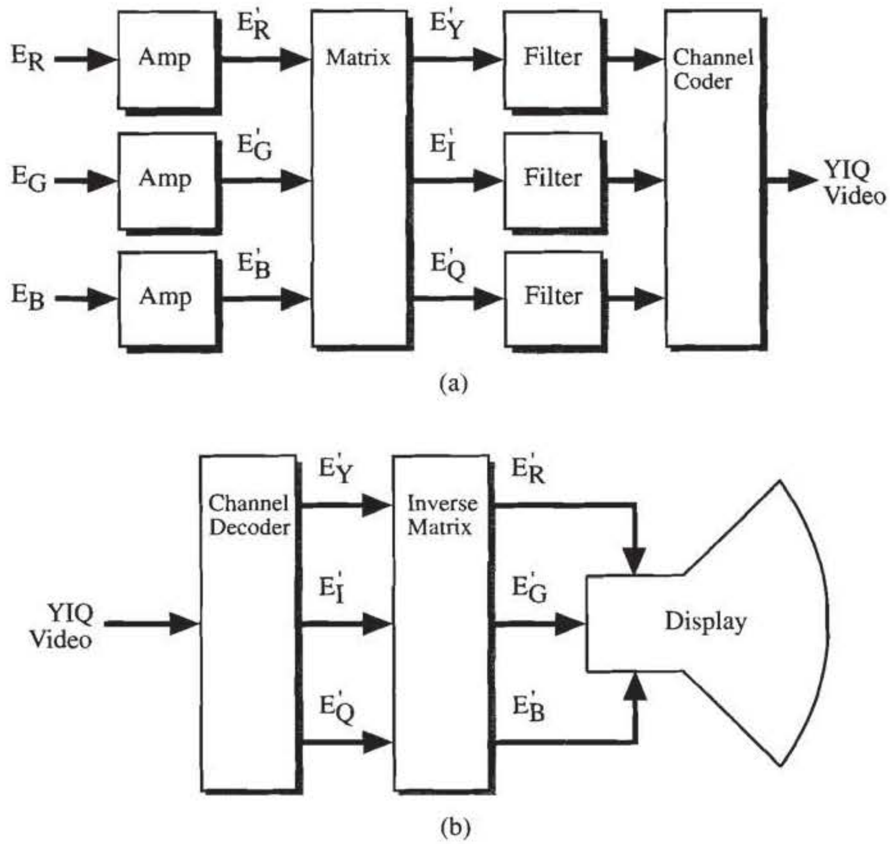

Figure 4. NTSC System (a) coder, (b) decoder

Because the RGB display at the sink is a power law device (light out is proportional to voltage in, raised to the power gamma γ), a compensating non-linear correction is required somewhere in the system to produce an overall linear characteristic. To minimize the visibility of channel noise and to reduce overall system cost (a broadcasting system has many decoders or receivers but few coders or transmitters), gamma correction was placed in the coder. Non-linear amplifiers convert the red, green and blue camera signals E_R, E_G, E_B into gamma-corrected signals E'_R, E'_G, E'_B.

The gamma-corrected signals are then matrixed to a luminance E'_Y and two chrominance signals E'_I and E'_Q (the inverse matrix is located at the decoder), which are filtered and then sent to the channel coder. The channel coder uses band sharing to fit the luminance E'_Y and two chrominance signals E'_I and E'_Q within the portion of the 6 MHz band allocated to the video signal; for details of the channel coder, see McIlwain and Dean (1956) for example.

The NTSC specifies a total systems solution based on 1950-era technology and user requirements. The situation has changed since then, and most of the design choices would be made differently today. In fact, these choices are being actively reexamined by the various systems proposed for HDTV (High Definition TV). One thing that hasn't changed is using luminance and chrominance color space for transmission and interchange.

The NTSC system assumes a standard RGB or display phosphors, although displays now use different, more efficient phosphors. Changing the decoder matrix can partly compensate for the color differences caused by the use of nonstandard display phosphors. Figure 5 plots standard phosphors for the European television system (Sproson, 1983), for studio monitors (SMPTE, 1986), and for HDTV (CCIR, 1990b) on a u'-v' chromaticity diagram.

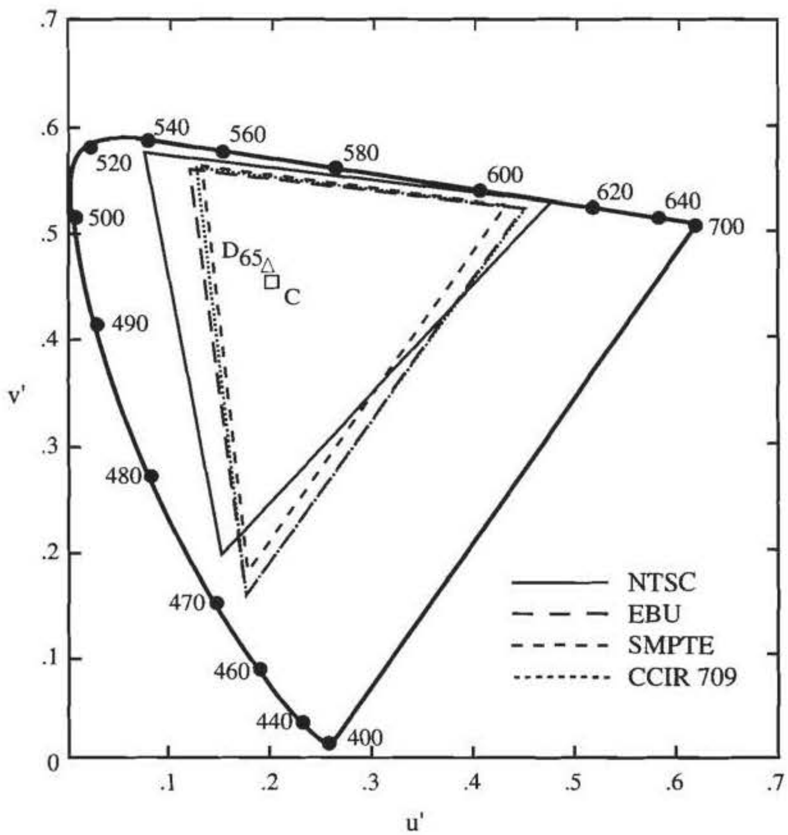

Figure 5. Standard Phosphors plotted on u'-v' chromaticity diagram

The evolution of phosphors has been accompanied by the development of new luma-chroma (generalized luminance-chrominance) spaces, designed to satisfy different requirements than the NTSC-YIQ space. Figure 6 shows the RGB phosphor sets and the luma-chroma spaces derived from them.

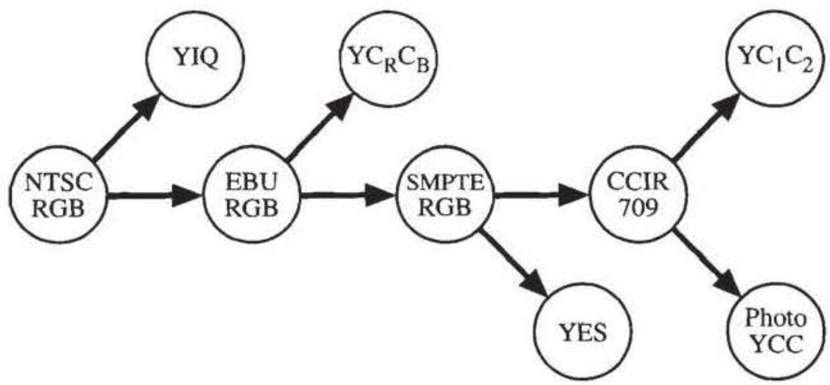

Figure 6. Evolution of RGB Phosphors and Luma-Chroma Spaces

YC_RC_B is used in European television broadcasting (CCIR, 1990a); YES is a linear space designed for color document interchange (Xerox, 1988); YC_1C_2 is recommended for HDTV systems (CCIR, 1990b); and PhotoYCC is the metric for the color image data contained on PhotoCDs (Eastman Kodak, 1991). With all these luma-chroma spaces, the bandwidth or spatial resolution of the chroma components can be reduced with little apparent loss in quality.

These luma-chroma spaces fit the following model:

$$\begin{bmatrix} L \\ C_1 \\ C_2 \end{bmatrix} = \underline{A} \begin{bmatrix} f_R(R) \\ f_G(G) \\ f_B(B) \end{bmatrix}$$

where the luma L and chroma coordinates C_1 and C_2 are a linear transform of functions f_R, f_G, f_B of R, G and B. For example, if f_R, f_G, f_B are non-linear gamma-correcting functions, then this model describes the NTSC coder. If XYZ tristimulus values are used in these equations in place of RGB, then the model applies to CIELab coordinates. Because CIELabspace has separate achromatic (L*) and chromatic (a* and b*) coordinates, it is also an example of a luma-chroma color space.

Digital Interchange Standards

Because of its storage, computational and programmable capabilities, a digital color system offers possibilities not found in an analog color system. The NTSC system is a hardwired system, using a fixed color space and fixed parameters for the transmission of scanned color data. In digital systems, the trend has been towards more flexible, programmable standards for interchange.

For example, the Xerox Color Encoding Standard was designed for the digital interchange of electronic documents (Xerox, 1988). It defines a small set of color spaces but with variable parameters that allow the user or source to select the dynamic range, number of bits, and tone scale compression of the color coordinates.

The ODA (Office Document Architecture) document interchange standard has even more flexibility, allowing the source to define both the color space as well as its parameters (ISO, 1991a). An ODA data stream can contain any color values as long as they can be transformed to the XYZ values of the CIE 1931 Standard Observer by the combination of matrices and lookup tables shown in Figure 7, or by a multidimensional lookup table. ODA also supports the use of CIELab and CIELUV, but only for images and not other kinds of color data.

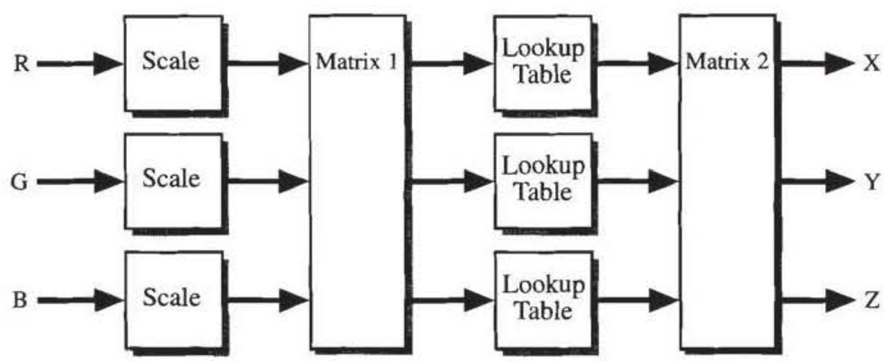

Figure 7. Office Document Architecture (ODA) Color Model

The block diagram in Figure 7 describes the color coder in an ODA system: the source values are labelled RGB and the interchange values are XYZ. In an ODA system, the parameters of all the coder elements can be set by the user; PostScript Level 2 (Adobe, 1990) and SPDL (Standard Page Description Language) (ISO, 1991b) have taken the same approach, but use procedures instead of lookup tables. Because the parameters of the coder are user specified, an ODA data stream must contain the coding parameters as well as the source color values.

In effect, ODA uses a "virtual" coder: no transform of the source values actually occurs at the coder. A possibly lossy transform from source to channel values is avoided by deferring the coding of source values to the decoder, where the coding and decoding can be combined into a single transformation from source to sink values. This approach also emphasizes backward compatibility, since existing color data can be inserted into an ODA data stream without modification, along with the parameters (i.e. the description of the coder) that would transform them to XYZ.

Interchange standards such as ODA, the Xerox Color Encoding Standard, PostScript Level 2 and SPDL use device-independent color, which means that the color values interchanged on the channel are based on the XYZ tristimulus values of the CIE 1931 Standard Observer. Because the source and the destination can use different media or viewing conditions, faithfully recreating the source tristimulus values at the destination is usually not desirable, and color interchange usually requires a transformation of the source color values.

Digital Color Coding

The discussion has so far mainly focussed on the color space used for interchange. In an analog system, the choice of color space affected bandwidth reduction. In a digital system, it affects the performance of the amplitude quantizing and spatial sampling operations, as well as the compression algorithms. The coding of digital color images is a subject on which books have been written (for example, Netravali and Haskell, 1988), but its treatment here must necessarily be brief and somewhat idiosyncratic.

Digital television standards provide for the 8-bit uniform quantization of both gamma-corrected RGB and luma-chroma components, with the chroma components sampled at half the rate of the RGB or luma components (CCIR, 1990a). Because gamma-corrected RGB or luma-chroma signals are visually more uniform than linear or camera RGB signals, they are more efficient color spaces for quantization, i.e. the same number of bits gives less visible quantizing noise, or the same noise visibility can be obtained with fewer bits.

One function of a digital coder is decorrelating the source signals, as correlated signals represent redundant and therefore superfluous information. In most natural scenes, RGB signals are highly correlated. A Karhunen-Loeve transform of RGB gives statistically uncorrelated components with the image energy optimally distributed between them (Pratt, 1971). Luma-chroma components have a similar energy distribution, and the luma-chroma transform in this sense approximates the Karhunen-Loeve transform.

There is one way in which RGB signals are more efficient than luma-chroma. In an additive system, all R, G, B combinations correspond to obtainable colors and no combinations or quantization levels are unused or wasted. The same is not true in a luma-chroma space. Figure 8a shows the chroma plane (in this case, the a*-b* plane in CIELab space) for a constant luma (L*) value. The shaded region is the slice through the color solid and contains all the colors obtainable at this L* value. Quantizing the color coordinates individually means defining levels for all values in $[a^*_{min}, a^*_{max}]$ and $[b^*_{min}, b^*_{max}]$. However, not all combinations of a* and b* values within these ranges correspond to obtainable colors. A measure of the efficiency is the ratio of the area of the shaded region to the area of the rectangle that encloses it.

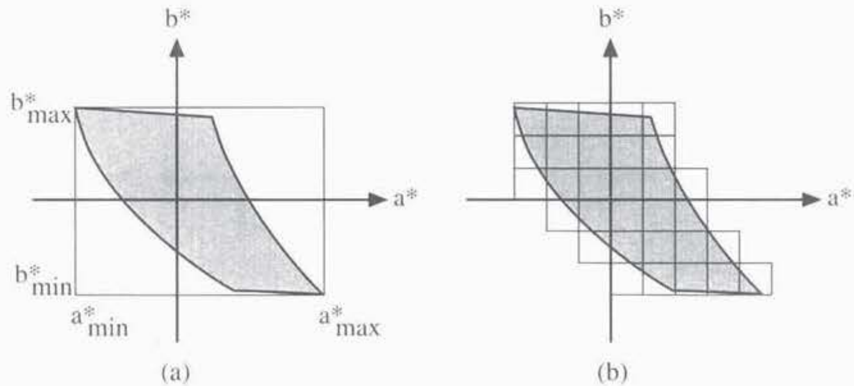

Figure 8. a*-b* Chroma Planes

Vector coding is a way to improve this ratio. Instead of coding the a* and b* components separately, they are coded in pairs or vectors (a*, b*), with codes provided only for the vectors that correspond to obtainable colors. A codebook or lookup table maps these vectors to codes. In the limit, there would be a table entry for each vector, which is equivalent to tiling the plane with unit squares. To keep the table to a reasonable size, the vectors input to the table could be $(a^*/2^m, b^*/2^n)$, which is equivalent to tiling the plane with $2^m \times 2^n$ rectangles, as shown in Figure 8b. Codes are only provided for the rectangles that intersect the shaded region of the plane or contain colors that actually occur in the image. Extending this to the 3-dimensional color solid would use (L*, a*, b*) vectors (Buckley, 1981).

In this example, vector coding addressed the correlation between color components. Vector coding can also address the correlation between spatial samples for each component, by assigning codes to blocks of samples rather than individual ones (Nasrabadi and King, 1988).

A key idea in coding is transforming the source signal to a form where the bits per sample or the number of samples can be reduced without a significant loss in picture quality. The luma-chroma transform was an example of such a transform; it produces chroma components that can be subsampled. Another example is the transform in the splitband coder shown in Figure 9 (Kretzmer, 1956).

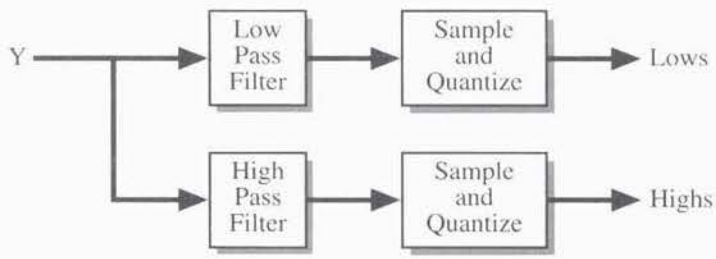

Figure 9. Splitband Coder

This coder splits the monochrome or luma signal into low and high frequency components. Fine detail and sharp edges are composed mainly of high frequency components, and the number of bits used to represent them can be reduced with little visible effect on picture quality. Doing the same for the low frequency components leads to visible quantization contours, but they can be subsampled.

Therefore, the lows are finely quantized but coarsely sampled, while the highs are finely sampled but coarsely quantized. This combination requires fewer bits than coding the luma signal directly with PCM. Kretzmer described the extension to multiple frequency bands, each coded by a different quantizer matched to the noise visibility properties of the band. This concept is now known as subband coding (Woods and O'Neil, 1986) and figures prominently in modern techniques for video coding.

Another version of the splitband coder is shown in Figure 10. The high frequency components are obtained by subtracting the low frequency components, or a decoded version of them, from the original signal. Cascading coders of this type gives a pyramidal coder (Burt and Adelson, 1983), where the lows output from one block are input to the next block. The high signal or residual from each block is coded and transmitted along with the lows signal or base image from the last block in the sequence.

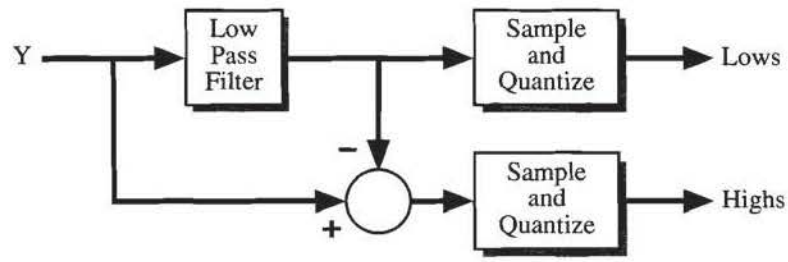

Figure 10. Alternate Implementation of Splitband Coder

How the highs or residuals are coded determines whether the coder is lossless or lossy. Lossless means that the output of the decoder and the input to the coder are identical; lossy means they are different. The goal in a lossy coder is to make the differences visually either undetectable or tolerable, i.e. "visually lossless." In general, the performance of a lossy coder is based on visibility criteria, and that of a lossless coder on signal statistics.

The original splitband coder is lossy because the coarse quantization of the highs cannot be reversed to recover the original values. The PhotoCD system uses a pyramidal coder for image storage (Eastman Kodak, 1991); it is a lossy coder because the residuals are quantized before they are Huffman coded, with high probability values assigned shorter codes than lower probability ones.

The splitband or two-channel coder resembles a mixed highs coder, and its extension to color is straightforward. The luma lows becomes a color lows signal, which is subsampled, while the luma highs are coarsely quantized. A coder of this type can compress a 24 bit per pixel color image into 2-3 bits per pixel, with only a slight loss in picture quality. Most of the 2 to 3 bits in the compressed image are devoted to the highs; the subsampled color lows require less than a bit per pixel (Schreiber and Buckley, 1981).

In both television and digital imaging systems, the cost in channel capacity of adding color is low. (The cost of color is borne by the coder and decoder, and by the source that generates it and the display that shows it.) A luma image can be turned into a color image by adding two low resolution chroma components. Schreiber (1991) has pointed out that adding color is actually a way of improving channel utilization, and that slightly reducing the luma resolution to make room for color almost always increases picture quality.

Splitting a signal into more and finer spatial frequency components eventually leads to the idea of a transform coder, shown in Figure 11 (Pratt, 1971). The filter bank of a splitband coder is replaced by a spatial transform that converts luma-chroma signals from the spatial domain (x, y) to the frequency domain (u, v). The quantizer operates on the signal in the spatial frequency domain.

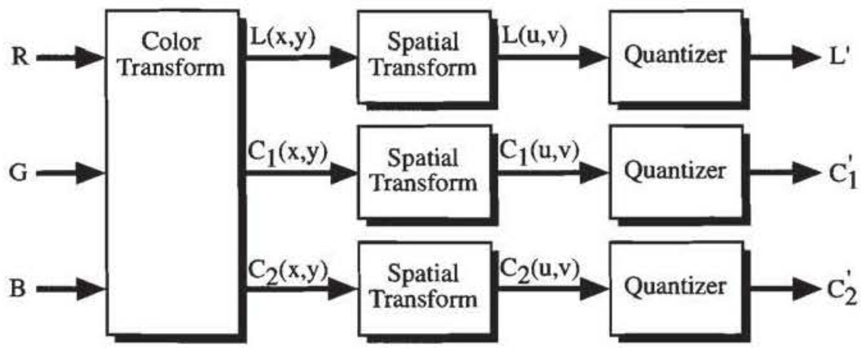

Figure 11. Transform Coder

The Joint Photographic Experts Group (JPEG) has recently developed a standard transform coder based on the discrete cosine transform (Wallace, 1992 and ISO, 1991c). While the standard defines offers several coding options, I will describe only the baseline coder, which all implementations must support and which is now available in VLSI from several vendors. The baseline coder for the luma channel is shown in Figure 12. The discrete cosine transform (DCT) converts an 8x8 block of 8-bit or 12-bit luma pixels L(x, y) into an 8x8 block of frequency domain samples L(u, v). By itself, this operation doesn't compress the image; it sets it up for the next stages which do.

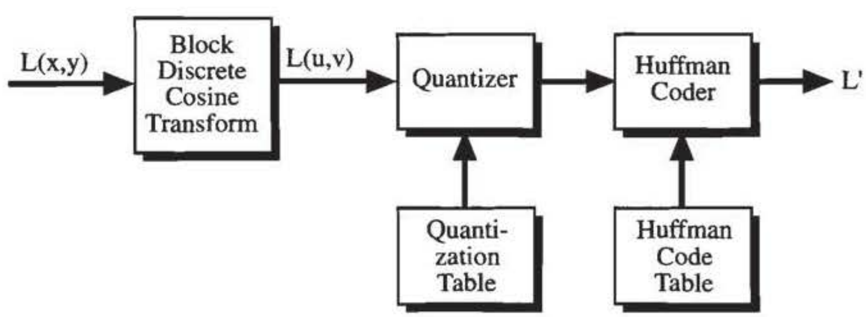

Figure 12. JPEG Baseline Coder for Luma Channel

The quantizer divides each frequency sample by a value from an 8x8 table. As a rule, these values increase with frequency so that the quantizer uses progressively fewer bits as the spatial frequency increases. The block data is then read out serially and the bit stream is Huffman coded. Because of the quantizer, the overall coder is lossy. The coder for the chroma is the same, except that the chroma data C(x,y) is subsampled before being transformed, and it uses different quantization and Huffman tables.

The JPEG algorithm can compress typical color images by factors of 10 or more with little loss in quality. Because Huffman coding is used, it is a variable bit rate algorithm, and different blocks or images can have different compression ratios. The compression ratio can also vary with the desired quality level.

In all this, the goal of the coder-decoder is producing a signal that is either identical or visually indistinguishable from the source signal. Because the systems we have been discussing are device independent, the source signals are tristimulus values, or a simple transform of them, such as luma-chroma. After briefly considering color coding in the human visual system, I will describe intermedia interchange where a decoded color signal that has the same tristimulus values as the source is not the goal.

Color Vision Models

Because the human visual system is the ultimate sink or destination in an image communications systems, its properties are used in the design of lossy image coders. A basic model for human color vision in shown in Figure 13. The three receptor responses R'G'B' are nonlinear functions of RGB, which are the integrals of the product of the receptor spectral sensitivity curves and the incident light.

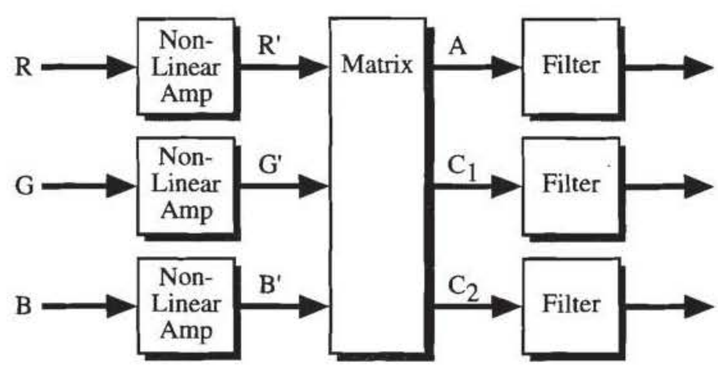

Figure 13. Human Color Vision Model

The R'G'B' signals are converted into an achromatic signal A and two color-difference or opponent-color signals C_1 and C_2. Buchsbaum and Gottschalk (1983) have pointed out that a Karhunen-Loeve transform of the receptor spectral sensitivity curves leads to an opponent-color representation, like the luma-chroma transform. Filters following this transform then shape the spatial frequency response of the three signals.

The configuration in Figure 13 can be traced as far back as Adams (1923). With XYZ inputs instead of RGB, the resulting achromatic and opponent-color signals were precursors of the L*, a*, and b* coordinates of CIELab space (Adams, 1942). This configuration describes the more recent models proposed by Hunt (1991) and Guth (1989). It also applies to the models with spatial filtering that were proposed by Faugeras (1979) and by Lloyd and Beaton (1990).

It is interesting to observe that the NTSC coder in Figure 4 resembles the model of Figure 13. Pearson (1975) noted that "it is not entirely unexpected that a system transmitting information to the eye should be found to be optimised by a similar form of [opponent] coding." Buchsbaum (1987) has noted the same resemblance. It is mainly superficial; as noted earlier, the NTSC coder is hardwired and fixed, while human color vision is marvelously adaptive.

Intermedia Interchange

The purpose of an image communications system is capturing an image at one point (the source) and recreating it at another (the sink). With the device-independent interchange standards described here, the connection between the source and sink uses XYZ tristimulus values, or an equivalent set of values such as luma-chroma, to represent the source image. Although the system uses tristimulus values, its goal is controlling appearance.

In a way that is easy to demonstrate but difficult to quantify, the context or viewing conditions associated with tristimulus values determine the mapping from them to appearance (Hunt, 1991, 1992). Because the source and sink can use different media with different viewing conditions, invariance of tristimulus values across the system is usually not the goal (Hunt, 1970). In these cases, the decoder will modify the source tristimulus values.

The color television broadcasting system has already had to deal with differences in viewing conditions between the source (the studio) and the sink (the home) environments. A typical situation is a studio scene illuminated by tungsten floods that is televised and viewed at home in a dim room on a receiver balanced for daylight. If the tristimulus values of the original studio scene were faithfully reproduced on the home receiver, the result would look yellowish and flat. It normally doesn't because the NTSC system compensates for the white point (or illuminant) and surround differences between the source and sink.

To compensate for different white points, the RGB signals at different stages in the system are balanced or equal for the local white illuminant. This way the RGB signals that are produced by the white reference illuminated by tungsten lamps in the studio are reproduced as daylight white on the home receiver. Buckley and Roetling (1992) have found that using device-independent RGB values for white point transformations gives better results than using CIELab or CIELUV values in a color copying system.

Compensating for the dim viewing surround of the home receiver requires increasing the contrast of the decoded image (Bartleson and Breneman, 1967). It was noted earlier that a television coder uses gamma correction to compensate for display nonlinearities. If V is the decoded signal into the display and L' is the light out, then the display transfer characteristic can be written

$$L' = V^{\gamma_d}$$

where γ_d is the display gamma. The transfer characteristic of the coder is

$$V = L^{1/\gamma_c}$$

where V is the output of the coder and L is the source signal input to the coder. Combining these two equations gives the overall transfer characteristic

$$L' = L^{\gamma_d/\gamma_c}$$

If $\gamma_c = \gamma_d$, then the overall system would be linear. But if the coder gamma is set at 80% of a typical display gamma, the result is about a 25% increase in contrast, which offsets the decrease in contrast due to the dim surround. This mechanism makes assumptions about the relationship between the source and the sink. If they are incorrect, because the display is viewed with a brighter surround for example, then the observer would manually adjust the contrast to get the desired results.

In contrast to the NTSC system, the digital color interchange standards transmit information about the viewing conditions along with the color data. In the case of ODA, the source white point is given; in the case of PostScript Level 2 and SPDL, it's the source white point and optionally the black point. PostScript and SPDL use this and the corresponding data for the destination to perform a nonlinear mapping of values in a user-specified color space, such as XYZ or RGB. This mapping can be used to match appearance or dynamic range across the source and destination.

The Xerox Color Encoding Standard has introduced the notion of appearance hints for controlling several aspects of color reproduction, including the mapping from device-independent color values to appearance. For example, the "source environment" hint describes the context in which color was created at the source. The destination can use this hint in several ways, and can change the preferred way of interpreting it as new technology or understanding becomes available. Furthermore, different destinations can interpret hints different ways.

CONCLUSIONS

This has been a whirlwind tour of a broad area. To a large extent we are dealing with ideas and techniques that emerged when color television first appeared around 1950. But things are changing rapidly, and we can expect the next review of this area to deal with evolving color appearance models and emerging high definition color television technology.

LITERATURE CITED

Adams, E.Q.

1923. "A Theory of Color Vision," Psychol. Rev., vol. 36, pp. 56-76

1942 "X-Z Planes in the 1931 I.C.I. System of Colorimetry," J. Opt. Soc. Am., vol. 32, no. 3, pp. 168-173

Adobe

1990. "PostScript Language Reference Manual," (Addison Wesley, Reading, MA), 2nd. ed., 765 pp.

Bartleson, C.J. and Breneman, E.J.

1967. "Brightness Reproduction in the Photographic Process," Phot. Sci. Eng., vol. 11, no. 4, pp. 254-262

Bedford, A.V.

1950. "Mixed Highs in Color Television," Proc. I.R.E., vol. 38, no. 9, pp. 1003-1009

Buckley, R.R.

1981. "Digital Color Image Coding and the Geometry of Color Space," Ph.D. Thesis, MIT Dept. of Elect. Eng. and Comp. Science

Buckley, R.R. and Roetling, P.G.

1992. "White Point Transformations and Color Data Interchange," SID Digest of Technical Papers, Vol. XXIII, pp. 560-563

Burt, P.J. and Adelson, E.H.

1983. "The Laplacian Pyramid as a Compact Image Code," IEEE Trans. Commun., vol. COM-31, no. 4, pp. 532-540

Buchsbaum, G.

1987. "Color Signal Coding: Color Vision and Color Television," Color Res. Appl., vol. 12, no. 5, pp. 266-269

Buchsbaum, G. and Gottschalk, A.

1983. "Trichromacy, opponent colours coding and optimum information transmission in the retina," Proc. R. Soc. Lond. B, vol. 220, pp. 89-113

CCIR, International Radio Consultative Committee

1990a. "Encoding Parameters of Digital Television for Studios," Rec. 601-2

1990b. "Basic Parameter Values for the HDTV Standard," Rec. 709

Eastman Kodak

1991. "Kodak Photo CD System: A Planning Guide for Developers" (Kodak Information Center, Rochester), 25 pp.

Faugeras, O.D.
1979. "Digital Color Image Processing Within the Framework of a Human Visual Model," IEEE Trans. Acoust., Speech, Signal Proc., vol. ASSP-27, no. 4, pp. 380-393

Guth, S.L.
1989. "Unified Model for human color perception and visual adaptation," Human Vision, Visual Processing, and Digital Display, SPIE vol. 1077, pp. 370-390

Hunt, R.W.G.
1970. "Objectives in Colour Reproduction," J. Photogr. Sci., vol. 18, pp. 205-215
1991. "Revised Colour-Appearance Model for Related and Unrelated Colours," Color Res. Appl., vol. 16, no. 3, pp. 146-165
1992. "Appearance Modelling," these proceedings

IRE, Institute of Radio Engineers
1954. Special issue on NTSC, Proc. I.R.E., vol. 42, no. 1

ISO, International Standards Organization
1991a. ISO 8613, Office Document Architecture and Interchange Format, Amendment 2 (Colour)
1991b. ISO 10180, Standard Page Description Language, Draft International Standard
1991c. ISO 10918, Digital Compression and Coding of Continuous-tone Still Images, Draft International Standard

Kretzmer, E.R.
1956. "Reduced-Alphabet Representation of Television Signals," IRE Nat. Conv. Rec., vol. 4, no. 4, pp. 140-147

Lloyd, C.J.C. and Beaton, R.J.
1990. "Design of a spatial-chromatic human vision model for evaluating full-color display systems," Human Vision and Electronic Imaging Models, Methods and Applications, SPIE vol. 1249, pp. 23-37

Loughlin, B.D.
1951. "Recent Improvements in Band-Shared Simultaneous Color-Television Systems," Proc. I.R.E., vol. 39, no. 10, pp. 1264-1279

McIlwain, K. and Dean, C.E.
1956. "Principles of Color Television" (John Wiley & Sons, New York), 595 pp.

Nasrabadi, N.M. and King, R.A.

1988. "Image Coding Using Vector Quantization: A Review," IEEE Trans. Commun., vol. 36, no. 8, pp. 957-971

Netravali, A.N. and Haskell, B.G.
1988. "Digital Pictures" (Plenum Press, New York), 586 pp.

Pearson, D.E.
1975. "Transmission and Display of Pictorial Information" (Halstead Press), 225 pp.

Pratt, W.K.
1971. "Spatial Transform Coding of Color Images," IEEE Trans. Commun. Tech., vol. COM-19, no. 6, pp. 980-992

Schreiber, W.F.
1991. "Fundamentals of Electronic Imaging Systems" (Springer Verlag, Berlin), 304 pp.

Schreiber, W.F. and Buckley, R.R.
1981. "A Two-Channel Picture Coding System: II - Adaptive Companding and Color Coding," IEEE Trans. Commun., vol. COM-29, no. 12, pp. 1849-1858

SMPTE, Society of Motion Picture and Television Engineers
1986. "Color Television Colorimetry," Recommended Practice 145

Sproson, W.N.
1983. "Colour Science in Television and Display Systems" (Adam Hilger, Bristol), 222 pp.

Wallace, G.K.
1992. "The JPEG Still Compression Standard," IEEE Trans. Consumer Electr., vol. 38, no. 1, pp. xviii-xxxiv

Woods, J.W. and O'Neil, S.D.
1986. "Subband Coding of Images," IEEE Trans. Acoust., Speech, Signal Proc., vol. ASSP-34, no. 5, pp. 1278-1288

Xerox
1988. "Color Encoding Standard" (Xerox Systems Institute, Palo Alto), 64 pp.