Experience with the new color facsimile standard

Giordano Beretta, Konstantinos Konstantinides, Ho John Lee,
Andrew Mutz

Hewlett-Packard Laboratories, 1501 Page Mill Road, Palo Alto, CA 94304

Abstract

An international standard for the exchange of continuous-tone color and gray-scale images has been approved under the auspices of the International Telecommunications Union – Telecommunications Sector (ITU-T, formerly known as CCITT). After reviewing the highlights of the new capabilities, we report on our experience in implementing this standard and testing the protocol. We believe the quality of images transmitted with color facsimile is sufficient for many communication tasks of people working in the various color areas.

1 Overview of the color-facsimile standard

Research and development of color facsimile can be dated prior to World War II [5]. This effort has resulted in expensive and proprietary systems for special applications, such as the transmission of color photographs by news delivery services. Three breakthroughs have occurred that will soon make color facsimile as ubiquitous as plain paper facsimile is now.

The first break-through has occurred in electronic imaging. New lossy data compression algorithms that discard image data that cannot be perceived allow considerable reduction of the image data size in bytes. This size reduction makes it viable to transmit images over the public-switched telephone network (PSTN) using ordinary modems.

The second breakthrough is in the cost of hardware. High performance processors, color scanners, and color printers have become so economical that they can be assembled into a color facsimile machine with a very accessible purchase price.

Last but not least, an international standard has been sanctioned [3, 4]. This means that users can universally exchange color images with remote users without bothering about the manufacturer or type of the facsimile stations.

1.1 Selection of a color space

In a study proposal introduced 1990 by the Nippon Telegraph and Telephone Company of Japan (NTT), business color images have been classified into four categories:

- full color (color photographs)
- multi-color (color charts and graphs)
- bi-color (documents marked up with red ink)
- mixed color (combinations of above documents, such as color pages of magazines)

Seventeen color spaces have been considered. They have included device spaces likes CMYK and YIQ, colorimetric spaces like the CIE color spaces, and color order systems such as the Munsell Renotation System.

These color space candidates have been evaluated using criteria relevant for the intended use of color facsimile mentioned above. Examples are quantization error, compatibility with compression algorithms, and color stability with white point change.

This evaluation has resulted in the selection of the CIE 1976 ($L^*a^*b^*$) [1] color space (CIELAB) as the mandatory color space for color facsimile. The standard also specifies CIE Standard Illuminant D_{50} as the default, with tristimulus values $X_0 = 96.422$, $Y_0 = 100.000$, $Z_0 = 82.521$ for the standard white, implicitly basing the color facsimile standard on the 1931 Standard Colorimetric Observer.

1.2 Color representation

The CIE system of colorimetry provides methods of predicting the magnitude of perceived color difference when two colors do not match. The commercial application of color requires not only specifying a desired color, but also specifying the deviation that can be tolerated in its reproduction. In fact, the CIE recommends the use of approximately uniform color spaces such as CIELAB to promote uniform practice in color difference evaluations.

The color facsimile standard does not specify a color difference formula and has no means to communicate color tolerances. For example, a sender cannot tune its tolerances depending on the tolerances of the receiving machine’s print engine. The tolerances are set by the transmitting station when the color gamut range is selected and the quantization errors in the encoded chrominance channels are set. We can speculate that color fidelity will vary widely, depending on the skills of the engineers designing a particular color facsimile machine and any trade-off they decide to embrace to pursue a particular market.

Special care has been necessary for the selection of an appropriate gamut range. Because the standard prescribes the allocation of 8 bits/pel/component (or optionally 12), a trade-off is necessary between the gamut size and the artifacts introduced by quantization into a limited data volume. It has been determined experimentally that with a gamut range of $a^* = [-80, 80]$, $b^* = [-80, 120]$, quantization artifacts under aggressive compression are relatively unobjectionable [7].

The standard prescribes a default gamut range of $a^* = [-85, 85]$, $b^* = [-75, 125]$. For those cases in which *a priori* knowledge about the gamut range is available and high color fidelity is necessary, the standard allows the specification of a custom gamut range. This custom gamut range can be changed on a *per* page basis.

When the default gamut range is used, the following representations of encoding CIELAB values as eight bit integers L, a, b result:

$$L = \frac{255}{100}L^*$$

$$a = \frac{255}{170}a^* + 128$$

$$b = \frac{255}{200}b^* + 96$$

These values are rounded to the nearest integer and clamped to the range [0, 255]. There is a similar representation scheme for optionally representing the values as 12 bit quantities instead of usual 8 bits.

1.3 Selection of a data compression method

For the image compression method, baseline JPEG (Joint Photographic Experts Group) [2] has been adopted. As shown in Fig. 1, this method consists of three stages. First, a sequential discrete cosine transform (DCT) is applied, which is an orthogonal and separable transform that allows near-optimum energy compaction and for which a number of fast algorithms with low computational complexity have been developed.

In a second stage the data is quantized based on the discrete quantization table (DQT). This stage is lossy and implementors design the DQT so that no visible artifacts are introduced in the image. The compression ratio of a file can be increased by setting a so-called *q-factor* or *scaling factor*, which is essentially a uniform multiplicative parameter that is applied to the quantization tables. Even when the tables are carefully designed to be perceptually lossless, a large *q*-factor will introduce artifacts, such as blockiness in areas of constant color or ringing on text characters.

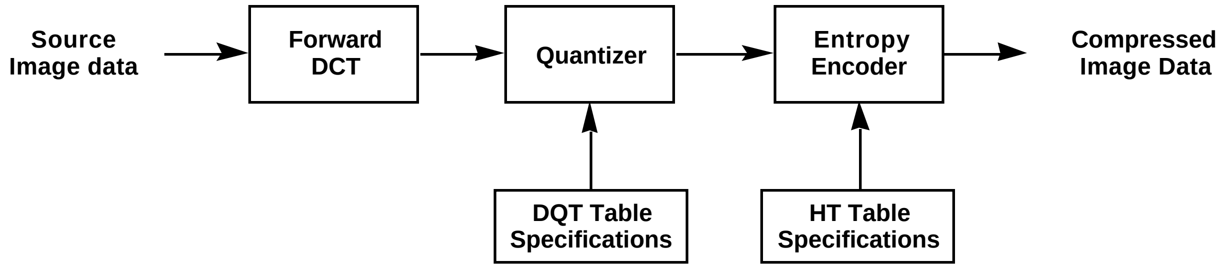

Figure 1. DCT-based encoder, simplified diagram. The lightness channels and the two chromatic channels are each run through the encoder. Typically, the tables are different for the various channels.(After reference 2, Fig. 4.)

The final step is a lossless Huffman encoder, which eliminates the entropy in the image file. The Huffman table (HT) controls the effectiveness of the lossless compression. At the receiving end, the inverse transforms are applied in reverse order, creating a rendition of the original image.

It has been established [6] that the luminance acuity is 32 to 33 cycles/degree, the red–green chromatic acuity is 11 to 12 cycles/degree, and the blue–yellow chromatic acuity is around 11 cycles/degree. Therefore, by default the color facsimile standards prescribes subsampling the image by 4:1:1 in the CIELAB space. For critical applications, it is possible to override this spatial subsampling in the chrominance channels.

1.4 The color processing pipeline

A Group 3 color facsimile machine essentially consists of an RGB scanner, a color space mapping module, a JPEG coder, a module performing the facsimile protocol, a modem, and a CMYK printer. Fig. 2 illustrates the color pipeline.

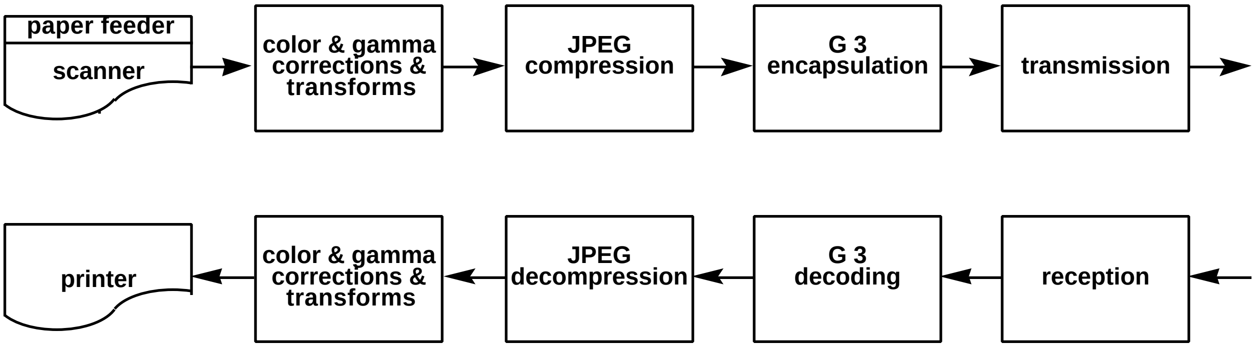

Figure 2. Color facsimile pipeline. The RGB bitmap from a scanner is converted to CIELAB color space and the chrominance channels are spatially subsampled. The then compressed using the JPEG method and encapsulated for the conventional Grou simile protocol and transmitted over an ordinary phone line. At the receiving end, th operations are applied.

The process begins with the translation of a color, 24-bits per pixel, RGB image into the CIELAB color space. At the baseline resolution of 200 dpi, an A size color image requires approximately 11.22M bytes of storage. Next, the chrominance components of the image are subsampled by a factor of two on both the horizontal and vertical coordinates. Subsampling decreases the size of the original image by a factor of two. Before transmission, the image is compressed using the baseline lossy JPEG algorithm. The compressed image is finally transmitted using the Group 3 facsimile protocol. On the receiver side, after Group 3 decoding, the image is decompressed and printed on a color or gray-scale printer.

The use of a device independent color space, a well defined gamut and the specification of the viewing conditions allow for a colorimetrically correct rendition of the image at the receiving color facsimile station. For artists, scientists, industrialists, educators, and the many others using color professionally, the characterization of the process is important also to know the limits of the system. This knowledge is important because the image is received remotely and by the very nature of the facsimile medium the sender can not compare visually the reproduction with the original.

2 Adding color to the facsimile protocol

To explore the issues for a practical realization we have assembled a system from a number existing components. This approach has allowed us to focus on the system issues and to ignore implementation details. After a short implementation time we were able to successfully verify the protocol with an external independent implementation.

The first step in augmenting an existing facsimile implementation for continuous-tone and gray-scale images is to provide for the new bits in the handshaking frames during the pre-message procedure. In this phase the receiving machine advertises its capabilities to the sending machine, which compares them with its own capabilities and selects the best common set.

The pre-message procedure allows users to communicate through a wide variety of facsimile machines. If the sender is a color facsimile system, the following receiver capabilities have to be evaluated:

1. Is the error correction mode (ECM) enabled?
2. Is the resolution at least 200×200 pels/25.4 mm?
3. Is JPEG coding, gray-scale mode supported?
4. Is JPEG coding, full color mode supported?

Any Group 3 facsimile machine supports the conventional modified Huffman coding that is used in all existing machines. If any of the above questions is negative, the sender falls back to this mode, ensuring full backwards compatibility with existing equipment. As the third test suggests, the new color facsimile protocol also supports a high quality black & white continuous tone mode.

If necessary, this negotiation can be repeated for each page, allowing for the transmission of mixed documents. For example, if a multi-page document contains only one color page and one gray-scale page, while the other pages are binary (e.g., black text) only the two continuous-tone pages will be JPEG encoded and only the color page will be sent in color mode. This allows the economical use of telephone lines.

The other components required to assemble a color facsimile system are readily available: scanner & printer drivers, and a JPEG coder. The JPEG coders available in the public domain usually require the image to be represented either in RGB or in YUV coordinates. Adding the conversion to and from the CIELAB color space is straightforward.

3 Transmission experiments

For error free communication using a 9600 baud modem, experimental and simulation results showed that the effective facsimile transmission rate is approximately 64K bytes per minute. This includes negotiation start-up time and typical retransmission overhead. Using lossy JPEG and the default quantization tables, we have achieved perceptually lossless compression with compression ratios close to 20:1.

For such a compression ratio, and a transmission rate of 64K bytes per minute, color facsimile of an A size color image will require at least 8.6 minutes of communication time. This is certainly neither cost-efficient or practical. Therefore, higher compression ratios, at acceptable printing quality, are required for efficient color facsimile. For example, for a transmission time of less than three minutes, the minimum required compression ratio is 60:1.

In the JPEG standard, the compression ratio is controlled by two sets of tables: the quantization and the Huffman tables. The quantization tables affect the lossy-compression part. A common practice is to use the default JPEG quantization tables, but scaled with an appropriate q-factor as mentioned above. The Huffman tables control the lossless compression part. Most applications use the default tables. For high compression ratios and acceptable output, both of these tables will have to be optimized.

Prior work related to JPEG optimization addresses only the YUV or the RGB color spaces and focuses on perceptually lossless compression. Scanning or printing are seldom taken into consideration. By taking into consideration the complete color facsimile pipeline (scanning, color transformation, compression-decompression, and printing), we are able to improve significantly both the compressibility and the printing quality of color images. Using custom quantization tables we are able to improve the compression ratio by at least 20%, for similar visual quality.

Table 1 shows compression ratios for three color images using both the default (uniformly scaled) and custom quantization tables. All images were scanned at 200 dpi. The 4CP01 image is the color facsimile test image recommended by the ITU.

Image	Default table	Custom table
Real estate flier with photo	52:1	82:1
Book page with photos and text	53:1	63:1
4CP01 test chart	47:1	63:1

Table 1. Typical compression ratios. We compare the compression ratios obtained with the default tables used in most JPEG implementations and the custom tables we tweaked by hand for color facsimile. The latter allow for viable transmission times.

Table 2 shows the timing results we obtained on a Hewlett-Packard Vectra 486/66XM personal computer. For data communication we used a 9600 bps modem. The A-size color image was compressed down with a conservative ratio to approximately 200K bytes.

Subsystem	Time in Min:Sec
Peripherals (I/O, color transformations, halftoning)	1:35
Transmission (including handshaking)	3:32
Processing (JPEG decoding)	0:30

Table 2. Timing results. Typical execution times for the components of a color facsimile testbed. The image transmitted was the 4CP01 test chart.

4 An open problem

The design of good quantization tables is more an art than a science. Much of the literature on color perception deals with aperture colors. There is only limited discussion of the spatial perception of color other than the study of visual acuity and contrast sensitivity. Knowledge about spatial color perception would allow the design of better quantization tables, which would shorten the transmission times for color facsimile communication.

Concomitantly, such knowledge would also allow the prediction of color tolerances.

5 Conclusions

We have shown that the new amendments to the Group 3 standard to provide for color and gray-scale facsimile are sound and do not present any particularly difficult system issues. We have tested our implementation of the protocol with an independent realization and successfully exercised all aspects of the protocol.

With the default JPEG quantization and Huffman tables, the transmission time are not practical for commercial applications. We were able to improve the transmission times with custom tables, but more work is necessary in this area.

The color facsimile standard does not have provisions for color tolerances. The prediction of color tolerances in data compressed with the JPEG method is still an open problem.

References

1. CIE, *Colorimetry*, Second Edition, CIE Pub. No. 15.2, Central Bureau of the CIE, Vienna, 1986.
2. ISO/IEC JTC 1/SC 29/WG 10 N52, *Information technology —- Digital compression and coding of continuous-tone images* —, ISO/IEC IS 10918-1, 1992.
3. ITU (Q.5/8 Rapporteur), *Amendments to ITU-T Rec. T.30 for enabling continuous-tone colour and grayscale mode for Group 3*, COM 8-43-E, 1994.
4. ITU (Q.5/8 Rapporteur), *Amendments to ITU-T Rec. T.4 to enable continuous-tone colour and grayscale mode for Group 3*, COM 8-44-E, 1994.
5. M. Matsuki and D.T. Lee, “Developments in Color-Facsimile Standards: An Overview,” Proc. *SID 94 Digest*, 883-886, 1994.
6. Kathy T. Mullen, “The Contrast Sensitivity of Human Colour Vision to Red-Green and Blue-Yellow Chromatic Gratings,” in *J. Physiol.*, **359**, 381-400, 1985.
7. A.H. Mutz, “Effect of Gamut Range Upon Color Quantization During JPEG Coding,” Proc. *SID 94 Digest*, 891-893, 1994.