

Proposal for a Standard Default Color Space for the Internet—sRGB

*Matthew Anderson** , *Ricardo Motta†* , *Srinivasan Chandrasekar** , *Michael Stokes†*
**Microsoft, where, where*
†Hewlett-Packard, Palo Alto, California

Introduction

HP and Microsoft propose the addition of support for a standard color space, sRGB, within the Microsoft OS's, HP products and the Internet. The aim of this color space is to complement the current color management strategies by enabling a third method of handling color in the OS's and the Internet that utilizes a simple and robust device independent color definition that will provide good quality and backwards compatibility with minimum transmission and system overhead. Based on a colorimetric RGB color space well suited to Cathode Ray Tube (CRT) monitors, television, scanners, digital cameras, and printing systems, such a space can be supported with minimum cost to software and hardware vendors. Our intent here is to promote its adoption by showing the benefits of supporting a standard color space, the suitability of the standard color space we are proposing, describe some of the system issues and propose a methodology for its implementation on the Web.

Part 1—Definition of the sRGB Color Space

Recently the International Color Consortium has proposed and is positioned to provide breakthrough solutions to problems in communicating color in open systems. Yet the ICC profile format does not provide a complete solution for all situations.

Currently, the ICC has one means of tracking and ensuring that a color is correctly mapped from the input to the output color space. This is done by attaching a profile for the input color space to the image in question. This is appropriate for high end users. However, there are a broad range of users that do not require this level of color quality, a broad range of file formats that will never support color profile embedding, and a broad range of uses that discourage people from appending any extra data to their files. It is at this level that a common standard RGB color space becomes useful and necessary.

We expect applications and users that do not want the overhead of embedding profiles with documents or images to convert them to a common color space and store them in that format. Currently there is a plethora of RGB monitor color spaces filling this void. There is a need to merge the many standard and non-standard RGB monitor spaces into a single standard RGB color space. Such a standard could dramatically improve the color fidelity in the desktop environment. For example, if operating system vendors provide support for a standard RGB color space, the input and output device vendors that support this standard color space

could easily and confidently communicate color without further color management overhead in the most common situations. The two major factors of this RGB space are the colorimetric RGB definition and the gamma value of 2.2, along with a number of secondary details necessary to enable the clear and unambiguous communication of color.

Colorimetric RGB

The dichotomy between the device dependent (e.g. amounts of ink expressed in CMYK or digitized video voltages expressed in RGB) and device independent color spaces (such as CIELAB or CIEXYZ) has created a performance burden on the system because of the complexity of the transforms, and a reliability gap since the complexity and variety of the transforms make it hard to ensure that the system is properly configured.

To address these concerns and serve the needs of PC and Web based color imaging systems, we propose a colorimetric RGB specification that is based on the average performance of personal computer displays. This solution is supported by the following observations:

- Most computer monitors are similar in their key color characteristics—the phosphor chromaticities (primaries) and transfer function
- RGB spaces are native to displays, scanners and digital cameras, which are the devices with the highest performance constraints
- RGB spaces can be made device independent in a straightforward way. They can also describe color gamuts that are large enough for all but a small number of applications.

This combination of factors makes a colorimetric RGB space well suited for wide adoption since it can both describe the colors in a unambiguous way and be the native space for actual hardware devices. This, many readers will recognize, describes in a roundabout way what has been the practice in color television for some 45 years. This proven methodology provides highest performance where it is needed the most, the fast display of images in CRT monitors.

Gamma 2.2

For computer software and hardware designers the most significant aspect of the proposed space is the 2.2 gamma. Because gamma correction tends to be a topic surrounded by confusion, it is worthwhile spending a few paragraphs discussing it.

The non-linearity of the electro-optical radiation transfer function of CRTs is often expressed by the exponent gamma. This transfer function describes how much visible radiant energy (cd/m²) results from voltages applied to the CRT electron-gun. Because most of the other characteristics of CRT based computer monitors are linear (including DACs and video amplifiers) the resulting transfer function has the same gamma value determining its non-linearity.

$$I = A (k_1 D + k_2)^{\text{gamma}} \quad (1)$$

Where K1 and K2 are the system gain and offset, D is the normalized pixel value, A is the maximum luminance of the CRT and I is the resulting luminance.

The key point that we wish to convey here is that gamma is dependent only on the electron gun design, and the vast majority of monitors and TV sets in use today are based on designs that result, on average, in the value 2.2 for gamma. Most of the variation between computer monitors and between TV sets are due to the differences in system gain and offsets (K1 and K2), which are partially under control of the user in the form of contrast and brightness knobs. Unfortunately, the actual set-up is often not know, but the best CRT performance happens when the system offset puts the dark parts of the images at the CRT cut-off, i.e. the black (pixel value 0) parts of the CRT image are just about to emit light. Under these conditions equation 1 above becomes

$$I = AD^{\text{gamma}} \quad (2)$$

and the monitor has the widest-dynamic range. The simplified form in equation 2 is what is usually found in the computer literature.

The gamma value 2.2 was the standard for television encoding before the advent of color TVs and was formalized in 1953. It is well matched to the eye's own non-linearity and it helps minimize transmission noise in the dark areas. Because all television sets have to display content generated with this encoding, it was very important for all CRT designs to conform to it. Only recently has the computer monitor market become as large as the TV market. As a result, most computer monitors have a native 2.2 gamma, with widest variations being in the set-up and screen reflectivity (older and less expensive display can reflect up to 20% of the ambient light). These factors typically can not be characterized a priori since they might change in the course of the day (ambient light) and at the whim of the user (by modifying contrast and brightness) but in practice the process tends to be self-regulatory, with users looking for darker places to set their monitors and modifying the controls to re-establish the expected display appearance. Exhaustive testing carried on at Hewlett-Packard on VGA computer monitors from many brands has shown the average gamma to be indeed 2.2, with a standard deviation of about 0.2.

Two special circumstances will lead computer systems to systematically deviate from the 2.2 gamma response - color dithering for 16 color systems and system imposed gamma correction via look-up-tables (LUT).

The first of course was very common until a few years ago. Until about 1993 most Windows PCs were well described by a gamma of 1.8 because despite having 2.2 gamma display systems, the color were dithered into the 4 bit frame buffers, resulting in a flattening of the system transfer function. This happens because screen dithering mixes colors linearly in the eye, making it less dependent of the CRT non-linearity. Since currently most Windows PC support 16 or 24 bit color modes, 2.2 gamma is now the average.

The second systematic deviation happens when the graphics system in the computer hardware or software imposes its own gamma correction. This is done for a variety of reasons, but is usually an attempt to compromise between image display and graphics/image processing performance (most computer graphic rendering assumes linear radiation space, e.g., transparency operations, and so does image processing, e.g., scaling and filtering). The gamma correcting of image data can be described by applying an exponent (g²) to the image data as follows :

$$I = A (D/g^2)^{\text{gamma}} \quad (3)$$

and the apparent system transfer function becomes:

$$I = AD^{\text{gamma}/g^2} \quad (4)$$

For the Macintosh the apparent exponent (gamma/g²) is around 1.8 and for SGI workstations around 1.7.

In summary, there has been some concern with the choice a 2.2 gamma as opposed to a 1.8 gamma. We feel that there are many reasons to support this choice, including compatibility with a large legacy of images (Photo CD, many Unix workstations, PC's with 256+ colors and their desktop color schemes and icons, several ultra-large image collections, analog television and CCIR 601 images), it is also a better fit with Weber's fraction, it is compatible with numerous standards (TIFF/EP, EXIF, digital TV, HDTV, analog video, and PhotoCD), it is closer to native CRTs gamma and it consistent with a larger market of displays.

Video Camera Encoding and ITU-R BT.709 Compatibility

There has been some confusion in previous versions of this proposal with respect to compatibility with the ITU-R BT.709 standard and with respect to the 709 recommendation in general. The ITU-R BT.709 standard specifically describes the encoding transfer function for a video camera that when viewed on a "standard" monitor will produce excellent image quality. The implicit target of this encoding is a standard video monitor whose transfer function is *not* explicitly delineated. Instead a typical monitor setup is assumed. This paper attempts to explicitly describe a standard monitor setup that is compatible with the 709 encoding standard.

This is illustrated in Figures 1-3 below. Figure one is directly derived from the ITU-R BT.709 standard. This standard provides mathematical methods to transform from tristimulus values of the scene using a video camera into a reference monitor device space. Figure two expands the implicit step of these methods and shows the transformation between the original scene tristimulus values into the target monitor tristimulus values. Since these two viewing

conditions are different, an implicit compensation is made to account for these differences (i.e., flare and ambient luminance). In order to provide an independent monitor reference color space, the monitor compensation methods must be extracted from this confounded compensation. This is precisely the goal of the sRGB color space. Figure three illustrates both the sRGB color space and the extraction of the monitor only specifications implicit within the ITU-R BT.709 standard. By producing such a monitor space, one can then transfer the ITU-R BT.709 encoded signals to other devices. By building on this system, the sRGB color space provides a monitor definition that can be used independently from the ITU-R BT.709 standard while maintain compatibility. This allows for the well-defined transfer of color information across the world wide web as described in the other section of this paper.

Another confusion often found when implementing the 709 specification is the actual differences between the transfer equations and the references to a gamma of 2.22. The actual 709 encoding transfer function can be fit with a power function closer to 1.95 than 2.22. This is due to the large offset of 0.099 in the transfer function equations.

The 709 transfer function in combination with its target monitor is attempting to achieve a system “gamma” of 1.13, assuming a monitor gamma of 2.5 and a encoding transfer function of 2.22 as shown in the equation below.

$$\text{gamma}_{\text{system}} = \text{gamma}_{\text{monitor}} / \text{gamma}_{\text{encoding}} \quad (5)$$

$$1.13 = 2.5 / 2.22 \quad (6)$$

Where the system gamma is 1.13, the monitor gamma is 2.5 and the encoding gamma is 2.22. Using the actual power function fit value for the 709 transfer function of 1.95 and maintaining the system gamma of 1.13, we can solve for the ideal target monitor gamma of 2.2. This is consistent with the gamma value proposed in this paper. This video compatibility also allows one to use one of the two current gamma paths that valid for the Apple Macintosh operating system (1.8 and 2.2).

$$1.13 = 2.2 / 1.95 \quad (7)$$

Viewing Environment Compensation

The reason that a system gamma of 1.13 is used instead of 1.0 is to compensate for the viewing environment conditions, including ambient illumination and flare. Historically, system gammas of 1.5 have been used for viewing projected slides in a dark room and system gammas of 1.25 have been used for viewing monitors in a very dim room. This very dim room value of 1.25 has been used extensively in television systems and assumes a ambient luminance level of approximately 15 lux. The current proposal assumes an encoding ambient luminance level of 64 lux of a more representative dim room in viewing computer generated imagery and therefore assumes a system gamma of 1.13, consistent with the 709 standard described above.

If the viewing condition is different from the standard, then the decoding process must compensate. This can be done by modifying the gamma values in equation 1.2 below by the appropriate factor. If one does modify the gamma values in equation 1.2 below, extreme care must be taken to avoid quantization errors when working with 24 bit images and high viewing flare levels.

Alpha Channel Masking and Computer Graphics Compatibility

Another concern that has been expressed about encoding using a gamma of 2.2 is the use of alpha masking. Typically computer graphics effects, including alpha masking, operate in an optical intensity environment as opposed to a visually uniform one. This mandates using a linear gamma of 1.0 in most computer graphic operations which is obviously incompatible with the visually uniform encoding using a gamma of 2.2. In an 24 bit encoding scheme, encoding a mid-level gray using a gamma of 1.0 would result in a digital count of 46. Such an encoding scheme would create visually objectionable artifacts such as contours. Therefore we recommend that effects such as alpha masking be performed either prior to encoding or by decoding to a color resolution greater than 24 bits and then converting into linear intensity space.

Again, it is fundamental to realize that there are two common operational spaces discussed here; 1) the linear intensity space in which most optical and synthetic visualization operations work within and 2) a nearly visually perceptibly uniform space in which visual based operations, compression and many devices work within. This difference is vitally important to remember since both types are operations are required and thus interaction between these two types of spaces occur frequently.

A Single RGB Standard Color Space

For a single color space to achieve acceptance, it must be objective, that is, have a tightly-defined relationship with the CIE standards. We are fortunate to have obtained in April 1990 unanimous worldwide agreement on a calibrated nonlinear RGB space for HDTV production and program exchange: Rec. ITU-R BT.709. This recommendation specifies the encoding of real world scene tristimulus values into a standard monitor RGB color space assuming a dark viewing condition. HP and Microsoft suggest using these parameters as the basis for the sRGB color space but with a dim viewing condition which is closer to most typical viewing environments for computer displayed imagery. Unfortunately the ITU specification is rather vague on the defining the target monitor. This paper attempts to provide a clear and well defined target monitor for the Rec. ITU-R BT.709 camera encoding standard for a dim viewing environment.

To be more precise, the ITU standard defines the transformation of real world CIEXYZ tristimulus values into a target RGB monitor space. This is essentially a composite of two transformations; one from real world CIEXYZ tristimulus values into standard monitor CIEXYZ tristimulus values and two from these standard monitor CIEXYZ tristimulus values into standard monitor rgb values. The re-

sulting image is not an exact appearance match of the original scene, but instead is a preferred reproduction of the original scene that is consistent with the limitations of a monitor.

This sRGB recommendation essentially defines the second part of this transformation between the target RGB monitor space and the monitor CIEXYZ tristimulus values in a dim viewing environment. This is required to maintain a consistent monitor centric color reproduction process which is typical of the web and is consistent with recommendations of the International Color Consortium. This is also consistent with maintaining the preferred reproduction encoding of the ITU standard.

There are two parts to the proposed standard described in this paper; the viewing environment parameters with its dependencies on the human visual system and the standard device space colorimetric definitions and transformations. The viewing environment descriptions contain all the necessary information, when combined with most color appearance models, to provide conversions between the standard and target viewing environments. The colorimetric definitions provide the transforms necessary to convert between the sRGB color space and the CIEXYZ two degree observer color space.

sRGB Reference Viewing Environment

Reference viewing environments are defined for standard RGB in Table 1.

The sRGB reference viewing environment corresponds to conditions typical of monitor display viewing conditions.

The *luminance level* is representative of typical CRT display levels.

The chromaticities of the *illuminant white* are those of CIE D65.

The *image surround* is defined as “20%” of the maximum white luminance. This is close to a CIELAB L^* value of 50, while maintaining computational simplicity. The areas surrounding the image being viewed are similar in luminance and chrominance to the image itself. This surround condition would correspond, for example, to a reflection print displayed on a spectrally non-selective gray background of about twenty percent reflectance, where the print and the background are uniformly illuminated by the same light source.

This specification defines the concepts of an encoding viewing environment and a typical viewing environment. The encoding viewing environment is consistent with recent discussions within the ISO JTAG2 committee for graphic arts and photographic viewing conditions for photographic images on monitors. If possible to achieve, this is the recommended viewing environment for viewing photographic images on monitors. The typical viewing environment is representative of a typical office or home office viewing environment for personal computers.

For optimal results, we recommend using the encoding viewing environment when viewing sRGB encoded images. We also recognize that this is quite different from typical viewing environment.

While one would theoretically use the viewing conditions which represent the actual or typical viewing environment, if this is done with 24 bit images a significant loss in the quality of shadow detail results. This is due to encoding the typical viewing flare of approximately 7.5 percent into a 24 bit image as opposed to the encoding viewing flare of 1 percent. Therefore we recommend using the encoding viewing environment for most situations including when one’s viewing environment is consistent with the typical viewing environment and not the encoding viewing environment.

The *encoding ambient illuminance level* is intended to be representative of a dim viewing environment. Note that the illuminance is at least an order of magnitude lower than average outdoor levels and approximately one-third of the typical ambient illuminance level.

The chromaticities of the *encoding ambient white* are those of CIE D50.

Encoding viewing flare is specified to be 1.0% of the maximum white-luminance level.

The *typical ambient illuminance level* is intended to be representative of a typical office viewing environment. Note that the illuminance is at least an order of magnitude lower than average outdoor levels.

The chromaticities of the *typical ambient white* are those of CIE D50.

Typical Viewing flare is specified to be 5.0% of the maximum white-luminance level.

Table 0.1. sRGB Viewing Environment Parameters

Condition	sRGB
Luminance level	80 cd/m ²
Illuminant White	$x = 0.3127,$ $y = 0.3291$ (D65)
Image surround	20% reflectance
Encoding Ambient Illuminance Level	64 lux
Encoding Ambient White Point	$x = 0.3457,$ $y = 0.3585$ (D50)
Encoding Viewing Flare	1.0%
Typical Ambient Illuminance Level	200 lux
Typical Ambient White Point	$x = 0.3457,$ $y = 0.3585$ (D50)
Typical Viewing Flare	5.0%

Colorimetric Definitions and Digital Encodings

sRGB in combination with the reference viewing environments can be defined from standard CIE colorimetric values through simple mathematical transformations.

CIE colorimetry provides the basis for sRGB encoding of the color. For the calculation of CIE colorimetric values, it is necessary to specify a viewing environment and a set of spectral sensitivities for a specific capture device. The definitions for RGB given in equations 1.1 to 1.3 are based on the colorspace’s respective viewing environment.

The CIE chromaticities for the red, green, and blue Rec. ITU-R BT.709 reference primaries, and for CIE Standard Illuminant D65, are given in Table 2.

TABLE 2. CIE Chromaticities for ITU-R BT.709 Reference Primaries and CIE Standard Illuminant

	Red	Green	Blue	D65
x	0.6400	0.3000	0.1500	0.3127
y	0.3300	0.6000	0.0600	0.3290
z	0.0300	0.1000	0.7900	0.3583

sRGB tristimulus values for the illuminated objects of the scene are simply linear combinations of the 1931 CIE XYZ values and these RGB tristimulus values can be computed using the following derived relationship:

$$\begin{bmatrix} R_{sRGB} \\ G_{sRGB} \\ B_{sRGB} \end{bmatrix} = \begin{bmatrix} 3.2410 & -1.5374 & -0.4986 \\ -0.9692 & 1.8760 & 0.0416 \\ 0.0556 & -0.2040 & 1.0570 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \quad (1.1)$$

In the RGB encoding process, negative sRGB tristimulus values, and sRGB tristimulus values greater than 1.00 are not typically retained. When encoding software cannot support this extended range, the luminance dynamic range and color gamut of RGB is limited to the tristimulus values between 0.0 and 1.0 by simple clipping. This gamut, however, is large enough to encompass most colors that can be displayed on CRT monitors.

The sRGB tristimulus values are next transformed to nonlinear sR'G'B' values as follows:

If $R_{sRGB}, G_{sRGB}, B_{sRGB} \leq 0.00304$

$$\begin{aligned} R'_{sRGB} &= 12.92 \times R_{sRGB} \\ G'_{sRGB} &= 12.92 \times G_{sRGB} \\ B'_{sRGB} &= 12.92 \times B_{sRGB} \end{aligned} \quad (1.2a)$$

else if $R_{sRGB}, G_{sRGB}, B_{sRGB} > 0.00304$

$$\begin{aligned} R'_{sRGB} &= 1.055 \times R_{sRGB}^{(1.0/2.4)} - 0.055 \\ G'_{sRGB} &= 1.055 \times G_{sRGB}^{(1.0/2.4)} - 0.055 \\ B'_{sRGB} &= 1.055 \times B_{sRGB}^{(1.0/2.4)} - 0.055 \end{aligned} \quad (1.2b)$$

The effect of the above equations is to closely fit a straightforward gamma 2.2 curve with an slight offset to allow for invertability in integer math. Therefore, we are maintaining consistency with the gamma 2.2 legacy images and the video industry as described previously.

Finally, the nonlinear sR'G'B' values are converted to digital code values. This conversion scales the above sR'G'B' values by using the equation below where WDC represents the white digital count and KDC represents the black digital count.

$$\begin{aligned} R_{8bit} &= ((WDC - KDC) \times R'_{sRGB}) + KDC \\ G_{8bit} &= ((WDC - KDC) \times G'_{sRGB}) + KDC \\ B_{8bit} &= ((WDC - KDC) \times B'_{sRGB}) + KDC \end{aligned} \quad (1.3)$$

This current specification proposes using a black digital count of 0 and a white digital count of 255 for 24-bit (8-bits/channel) encoding. The resulting RGB values are formed according to the following equations:

$$\begin{aligned} R_{8bit} &= ((255.0 - 0.0) \times R'_{sRGB}) + 0.0 \\ G_{8bit} &= ((255.0 - 0.0) \times G'_{sRGB}) + 0.0 \\ B_{8bit} &= ((255.0 - 0.0) \times B'_{sRGB}) + 0.0 \end{aligned} \quad (1.4)$$

This obviously can be simplified as shown below.

$$\begin{aligned} R_{8bit} &= 255.0 \times R'_{sRGB} \\ G_{8bit} &= 255.0 \times G'_{sRGB} \\ B_{8bit} &= 255.0 \times B'_{sRGB} \end{aligned} \quad (1.5)$$

The reverse relationship is defined as follows;

$$\begin{aligned} R'_{sRGB} &= R_{8bit} \div 255.0 \\ G'_{sRGB} &= G_{8bit} \div 255.0 \\ B'_{sRGB} &= B_{8bit} \div 255.0 \end{aligned} \quad (1.6)$$

If $R_{sRGB}, G_{sRGB}, B_{sRGB} \leq 0.03928$

$$\begin{aligned} R_{sRGB} &= R'_{sRGB} \div 12.92 \\ G_{sRGB} &= G'_{sRGB} \div 12.92 \\ B_{sRGB} &= B'_{sRGB} \div 12.92 \end{aligned} \quad (1.7a)$$

else if $R_{sRGB}, G_{sRGB}, B_{sRGB} > 0.03928$

$$\begin{aligned} R_{sRGB} &= \left[\frac{(R'_{sRGB} + 0.055)}{1.055} \right]^{2.4} \\ G_{sRGB} &= \left[\frac{(G'_{sRGB} + 0.055)}{1.055} \right]^{2.4} \\ B_{sRGB} &= \left[\frac{(B'_{sRGB} + 0.055)}{1.055} \right]^{2.4} \end{aligned} \quad (1.7b)$$

and

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.4127 & 0.3586 & 0.1808 \\ 0.2132 & 0.7172 & 0.0724 \\ 0.0195 & 0.1197 & 0.9517 \end{bmatrix} \begin{bmatrix} R_{sRGB} \\ G_{sRGB} \\ B_{sRGB} \end{bmatrix} \quad (1.8)$$

Digital broadcast television uses a black digital count of 16 and a white digital count of 235 in order to provide a larger encoded color gamut. We do not propose using this encoding at this time, due to the large legacy of images and applications using the previous black and white digital coding counts. However, it is vital to allow for a future revision to provide this capability.

Part 2—Implementation on the Web

Color Spaces

Definition: A color space is a model for representing color numerically in terms of three or more coordinates. e.g. The RGB color space represents colors in terms of the Red, Green and Blue coordinates.

For color to be reproduced in a predictable manner across different devices and materials, it has to be described in a way that is independent of the specific behavior of the mechanisms and materials used to produce it. For instance, color CRTs and color printers use very different mechanisms for producing color. To address this issue, current methods require that color be described using device independent color coordinates, which are translated into device dependent color coordinates for each device.

Color Management

Definition: Color management is a term that describes a technology that translates the colors of an object (images, graphics or text) from their current color space to the color space of the output devices like monitors, printers, Ö.

Traditionally, operating systems have supported color by declaring support for a particular color space, RGB in most cases. However, since RGB varies between devices, color was not reliably reproduced across different devices.

The high-end publishing market could not meet its needs with the traditional means of color support, so the various OS's added support for using International Color Consortium (ICC) profiles to characterize device dependent colors in a device independent way. They use the profiles of the input device that created an image and the output device that displayed the image and create a transform that moves the image from the input device's color space to the output device's color space. This resulted in very accurate color. However, it also involved the overhead of; transporting the input device's profile with the image and running the image through the transform.

HP and Microsoft propose an additional means of managing color that is optimized to meet the needs of most users without the overhead of carrying an ICC profile with the image: the addition to the OS and the Internet of support for a Standard Color Space. Since the image is in a known color space and the profile for that color space would ship

with the OS and browser, this enables the end users to enjoy the benefits of color management without the overhead of larger files. While it may be argued that profiles could buy slightly higher color accuracy, we believe that the benefits of using a standard color space far out-weigh the drawbacks for a wide range of users. The migration of devices to natively support the standard color space will further enhance the speed and quality of the user experience.

We are proposing the use of the color space, sRGB, that is consistent with but is a more tightly defined derivative of Rec. ITU-R BT.709 as the standard color space for the OS's and the Internet. This space obtained in April 1990 unanimous worldwide agreement as the calibrated nonlinear RGB space for HDTV production and program exchange.

Proposed Style Sheet Syntax for Specifying sRGB on the Internet

We propose that all page elements defined in the style sheets be assumed to be in the sRGB color space unless embedded ICC profiles indicate otherwise. We further propose the following style sheet syntax to specify the user rendering intents of page elements. These rendering intent values are defined to be consistent with the International Color Consortium rendering intent values (<http://www.icc.org>) where S is saturation, P is perceptual, A is absolute colorimetric and R is relative colorimetric.

Syntax

Value: SI|A|R

Initial: P

Applies to: all elements

Inherited: yes

Percentage values: N/A

Example:

```
BODY {rendering_intent: S}
IMG.perceptual_intent {rendering_intent: P}
#mypic001 {rendering_intent: A}
<IMG ID=myspic001 SRC="http://www.site.com/layout.mypic001.gif" >
<IMG CLASS= perceptual_intent SRC="http://www.site.com/layout.mypic002.gif" >
```

	Style Sheet Colors(sRGB)	HTML Page with no Color Space information	Re-purpose Data outside of Browser/HTML environment
Embedded Profile in Image	Color Space for Image determined by embedded profile.	Color Space for Image determined by embedded profile.	Color Space for Image determined by embedded profile.
Image file specifies sRGB	Color Space for Image is sRGB	Color Space for Image is sRGB	Color Space for Image is sRGB
Image has no Color space information.	Color Space for Image is sRGB	Color Space for image is sRGB.	Color Space for image is sRGB.
Text	Color Space for text is sRGB	Color Space for text is sRGB.	Color Space for text is sRGB.
Graphics	Color Space for Graphics is sRGB	Color Space for graphics is sRGB.	Color Space for graphics is sRGB.

Standard Color Space in Practice

Once page elements are converted to sRGB, the browser needs to interpret the color space correctly and use the OS color management to image the page. The following table summarizes how the browser handles color management in each of the possible scenarios.

Browsing Scenarios

The following cases describe what an end-user sees in the various scenarios:

1. Image not in sRGB, does not have an embedded ICC profile, and no monitor/output device ICC profile

This is the behavior before color management systems were added. Even though the image is assumed to be in sRGB color space, it is imaged (displayed, printed etc.) without translation to the device color space since the output profile is not available. The quality varies tremendously since output device characteristics differ greatly.

2. Image not in sRGB, does not have an embedded ICC profile, and system has a monitor/output device ICC profile

Since the image has no ICC profile, it is assumed to be in the sRGB color space. In this scenario, the resulting image will be consistent across devices; however it could be different from the original image.

3. Image in sRGB, and no monitor/output device ICC profile

In this scenario, the image has been run through a transform that consists of the input device ICC profile, and the sRGB ICC profile, or it was created using devices that conform to sRGB. However, since the system has no ICC profile for the output device, it will simply assume the image is in the device's color space. If all the images rendered on the output device are in sRGB, then they will at least be consistent with respect to each other.

4. Image in sRGB, and system has a monitor/output device ICC profile

In this scenario, the image has been run through a transform that consists of the input device ICC profile, and the sRGB ICC profile, or it was created using devices that conform to sRGB. Because the system has an ICC profile for the output device, the image can be converted to the output device's color space and imaged. The resulting image will be consistent across devices, and will be very close to the original in appearance.

5. Image in sRGB, and monitor/output device is sRGB compliant

In this scenario, the image has been run through a transform that consists of the input device ICC profile, and the sRGB ICC profile, or it was created using devices that conform to sRGB. As the output device

has been designed to conform with sRGB, and is associated with that ICC profile, a transform is not necessary for this case. The OS realizes that no transformation is required and simply images the image directly on the output device. This case is ideal since there is no color transformation at output time, and the image is more compact since there is no ICC profile embedded in it. The resulting image will be consistent across devices, and will be very close to the original in appearance.

6. Image not in sRGB, has an embedded ICC profile, and no monitor/output device ICC profile

This would be treated the same as the "Image not in sRGB, does not have an embedded ICC profile, and no monitor/output device ICC profile" scenario.

7. Image not in sRGB has an embedded ICC profile, and system has a monitor/output device ICC profile

This is the standard color management scenario. The two ICC profiles are combined to produce a transform that will transform the colors of the image into the output device's color space. The resulting image will be consistent across devices, and will be very close to the original in appearance.

Authoring Scenarios

The following scenarios describe how to get an image into the sRGB color space when creating it.

1. Image created on a device that has no ICC profiles and is not sRGB compliant

Display the image on a monitor that is sRGB compliant or that has an ICC profile. Edit the image until it looks good on the monitor. For monitors that are not sRGB compliant but have ICC profiles, depending on the capabilities of the application, embed the monitor's profile into the image, and use a tool to create a transform with the monitor's profile and the sRGB profile and run the image through the transform. If the image file format supports it, specify the image is in sRGB.

2. Image created on a device that has ICC profiles and is not sRGB compliant

Use a tool to create a transform with device's profile and the sRGB profile. Then run the image through the transform, specify the image is in sRGB if the image file format supports it.

3. Image created on a device that is sRGB compliant

Specify the image is in sRGB if the image file format supports this.

Suggestions to Benefit from sRGB

1. CRT manufacturers who build monitors in compliance with the sRGB specification will get faster display times for objects in this color space. (No transformation needs to occur.)

2. Scanner and digital camera manufacturers who optimize the color transforms and gamma correction for compatibility with sRGB will benefit for the same reason.

Palette Issues

There are several different scenarios to consider when dealing with palettized images and displays.

1. **Image does not have a color table (> 8bpp), and client monitor is not palettized**

The image is run through a color matching transform as described in the previous section, and the resulting 24bpp image is displayed on the monitor.

2. **Image has a color table (8bpp) and client monitor is not palettized**

The color table accompanying the image is run through a color matching, and the resulting color table is used with the image for display. The displayed image is very close to original image.

3. **Image does not have a color table (>8bpp) and client monitor is palettized.**

The software displaying the image (e.g., browser) should use the default palette that is defined in sRGB space, convert it into device color space by running it through a color matching transform, and use this palette to display the image. The resulting image gets dithered into the closest possible colors on the display. The assumption is that the monitor profile is created with the default palette selected.

4. **Image has a color table (8bpp) and was created using the default palette and client monitor is palettized**

The software displaying the image should follow the same steps as above. The resulting image is very close to the original image and unintentional dithering is eliminated. If the original image only had colors in the default palette, the final image doesn't have any dithering.

5. **Image has a color table (8bpp) and was created using an arbitrary palette and client monitor is palettized**

As the software displaying the image cannot distinguish this case from case 4 above, it would (and should) follow the exact same steps as 4. The resulting image is reasonably close to original image in most cases. This claim is made under the assumption that the default palette covers the device's color space rea-

sonably well and the source image isn't concentrated on one part of the spectrum. In the case where the colors of the source image are very localized, the final image could be quite heavily dithered. Authors publishing images on the Internet should avoid this case as far as possible because of this very problem.

Note that cases 3, 4, and 5 assume an industry standard default palette defined in sRGB color space that will be used by authoring and display software to handle 8bpp images. Microsoft and HP are working on the definition of such a palette.

Conclusion

We believe that the addition of standard color space, sRGB, support to the Internet and the OS is a complementary addition to the existing color management support that utilizes and expands the benefits and availability of color management to a broader range of users.

Call for Action

1. OS and authoring tools should utilize CRT calibration methods to ensure that images created locally on the monitor are properly translated into sRGB.
2. Authoring tools should enable the use of the OS color management tools to transform the incoming images into sRGB by combining the incoming device ICC profiles with the sRGB ICC profile. These images will be stored in the original file format (like GIF and JPEG), but with the correct colors.
3. Authoring tools should enable the user to view and edit text and graphics in sRGB.
4. Web page creators should publish content in sRGB.
5. Browser vendors should support the style sheet extensions proposed above and use the OS color management tools to ensure that colors are properly displayed.
6. Organizations that create and support file formats should ensure that the file formats encompass the ability to embed profiles and declare their color space.

Acknowledgements

The authors would like to sincerely thank Ed Giorgianni of Eastman Kodak Company, Bob Sobol of Hewlett-Packard Company and Charles Poynton, author of *A Technical Introduction to Digital Video* (John Wiley and Sons), for their many insightful comments and encouragement throughout this process.