Color Management Technology for Workstations

Color monitors are now a ubiquitous among workstations, and desktop scanners and printers are becoming increasingly affordable. However, mismatches in color handling mean that you cannot yet scan a color image, look at it on your screen, print it on a local color printer, and send it off to a prepress system, and expect the colors to remain true at each stage of the process. This situation is about to change: color management will soon be available on desktop workstations. [1]

CRT Displays

Because cathode ray tube (CRT) monitors have red, green and blue phosphors, computer graphics image data is conventionally represented in RGB color components. However, there is no universal, objective definition of what colors constitute “red”, “green”, “blue”, or even “white”. For RGB data that indicates saturated red — 100% red, 0% green, 0% blue — is the red intended to be scarlet, reddish-purple, or reddish-orange? In practice, the color reproduced depends upon the interpretations given to each of the primaries by a particular device. In a CRT monitor, the colors produced by the phosphors used in the manufacture of the CRT — the *phosphor chromaticities* — determine the colors of red, green and blue. The balance of power among the three electron beams — the *white point* — determines the color assigned to white. These parameters are different for different monitors, and color reproduction is not predictable without control of these parameters.

Desktop Scanners

The color interpretation of red, green and blue data from a scanner depends on the characteristics of the optical filters that are used to separate the wavelengths of visible light into color components. Current desktop scanners have notoriously poor color characteristics, and this causes scanners to fail to “see” colors the same way that we do.

Desktop Color Printers

When ink or other transparent colored material is deposited onto paper, the widest range of colors is produced when the colors of the inks are yellow, cyan and magenta. Current desktop color printers utilize CMY colorants, deposited

---

[1] This paper was presented at Sun Expo ‘92, Manchester, U.K. on Sepember 10, 1992.

by one of two dominant technologies: ink jet or dye sublimation. Ink jet printers deposit either a full droplet of ink, or no ink, at each of the device's pixels: reproduction of continuous-tone imagery is achieved by the technique of halftoning, which I will describe in a moment. Although the basic resolution of ink-jet printers can be quite high, in the range of 300 or 400 dpi, use of halftoning for picture reproduction reduces this resolution to about 50 lines per inch, comparable to the color quality of a newspaper.

A dye sublimation printer is capable of depositing a variable amount of colorant at each pixel, and consequently can produce a continuous-tone picture without suffering the impairments of halftoning. Dye sublimation printers are capable of near-photographic image quality, however this advantage comes at a premium price of about ten times that of a color ink-jet printer.

The exact color reproduced by a printer depends on the spectral characteristics of its pigments (inks) and its media, and upon the illumination under which the print is viewed. Characterizing the reproduction process of a printer is much complicated than for a CRT, because the overlap among the spectra the inks causes the color combinations to be nonlinear. For example, yellow ink is supposed to absorb only blue light, but practical yellow inks usually absorbs a significant amount of green as well. Achieving correct reproduction of mixtures that involve green requires knowledge of this interference.

Printing and Publishing

In traditional printing, there are fundamentally two ways to get a particular color onto the page. One way is to formulate a specific ink that produces the color that you require, and deposit that onto the page. This is known as solid color or spot color, and is appropriate for printing company logos onto stationery or printing a product's package in a color associated with the manufacturer. This approach is practical for problems involving a small number of colors that are known in advance, but is incapable of reproducing color pictures.

The second approach is to establish a set of several standard inks, and produce a range of colors using mixtures from that set. This approach is known as process color, and it almost always involves overprinting carefully-chosen amounts of four inks that are colored cyan, magenta, yellow and black (CMYK).

The fourth ink, black, is not necessary in theory. However, there are several practical reasons to use it. Colored ink is much more expensive than black ink, so performing gray component replacement (GCR) —on those areas of the image that would otherwise be overprinted cyan, yellow and magenta — saves a substantial amount of money. Replacement of three inks by one makes the printed page less wet with ink, which enables shorter drying time and faster printing speeds. Finally, registration errors tend to be less visible when the information that will be perceived as black-and-white is printed in a single pass. The black ink was historically called the key, and that is why it is abbreviated K.

Image data must necessarily at some point in the printing process be separated into CMYK components in order for an image to be printed using process color. Traditionally, color separation has been performed at the time that the original photographic imagery was scanned. Color separation involves very expensive machinery and a highly skilled operator. CMYK separation is very dependent on inks, paper and other factors of the printing processes.

Once an image data is in CMYK form it is in some sense married to a particular set of printing processes, inks and media. Once an image is in CMYK form, to move the printing job from one plant to another is difficult or impossible. Also, reliance on the CMYK color model makes it difficult to use the RGB devices, software and images that are ubiquitous in the computer domain.

Halftoning

The printing processes that achieve low cost reproduction lay down — at each position on the page — either a fixed amount of ink or no ink at all. The appearance of continuous-tone reproduction is obtained by laying down a fine pattern of dots of varying size. This process is called halftoning, and is in contrast to continuous tone processes such as photography that can deposit a variable amount of colorant at each point in the reproduced image.

When halftoning is implemented in the digital domain, the page is divided into a uniform non-rectangular array of cells. Each cell contains between a few dozen and a few hundred device pixels. To produce a progressively darker shade of gray, device pixels are “turned on” one-by-one, in a pattern that progressively enlarges the reproduced dot. The shape of the reproduced dot is determined by the arrangement of the cells and on the progression of device pixels that enlarges the dot as a function of the amount of ink required. Dot shape is a major determinant of image quality.

Obviously there is a tradeoff between a halftoning cell having a small number of pixels, which achieves good detail but a small number of shades of gray, and a cell having a large number of device pixels, which yields a large number of shades of gray but reproduces detail poorly. The pitch of the cells employed in practice ranges from about 50 lines per inch for the color comics to about 150 lpi for a high-quality magazine.

In color halftoning, each of the CMYK layers is associated with its own cell arrangement. The four screens typically have the same pitch but are reproduced at different, carefully-controlled angles.

Color mixing in a halftoned reproduction is very complex. When the dots are small, in the light tones of the image, the overlap among dots is minimal and the colors of the inks tend to mix additively. When the dots are large, in the dark areas of the picture, the overlap tends to be substantial and the colors mix “subtractively”. In the mid tones, the color mixtures are part additive and part subtractive. Also, in practice the amount of ink that is deposited on the page for each halftone dot is a nonlinear function of the size of the halftone dot that is formed on the printing plate. Compensation for this phenomenon of dot gain is necessary to achieve good quality reproduction.

Device-dependent Color

From the descriptions above, you can see that the reproduction of color in all of these systems — CRT displays, scanners, color printers and prepress systems — is highly device-dependent: if RGB values from one device are sent to another without consideration of different device characteristics, then different colors are likely to result. This is the typical situation today, except in highly-specialized systems that take proper account of device color reproduction characteristics.

Monitor calibrators are commonly used today, particularly in workstations employed in publication work. Calibrating a monitor gives a known color interpretation to RGB values used in an application. But monitor calibration alone does not help with input from scanners and output to printers. Scanner calibration software has recently been introduced that transforms a scanner’s RGB data into a different RGB system suitable for display on a particular calibrated monitor. But end-to-end color reproduction cannot be solved by piecemeal compensation and correction between each pair of systems that exchange color information: a systems approach is necessary.

Color Management Systems

Sun plans to remedy these problems by introducing color management system (CMS) technology. A CMS comprises software, resident on the workstation, that performs the mathematics of color transformation among devices. The CMS facilities are provided as a software library with an application program interface (API) that presents color transform capability to any application. The CMS maintains device characterization and calibration information for color devices in order to achieve device independence.

The CMS is structured as a high-level color management framework (CMF) that presents the interface to the graphics libraries and to the applications, and a number of low-level color management modules (CMMs), each of which has responsibility to handle the mathematical transformations for a certain class of devices. The Sun CMS will be provided with a built-in CMM to handle color transforms among Sun devices. The architecture is designed to accommodate additional third-party plug-in color manager modules (CMMs), which are expected to become available from various suppliers. Kodak has demonstrated the ColorSense™ system, and Electronics for Imaging (EFI) has demonstrated EFIcolor™. Traditional prepress vendors — such as Agfa, Crosfield, Dai-Nippon (Screen), Hell and Scitex — have systems that perform these tasks, and it is likely that these vendors will use standardized color management system interfaces to bring their high-end color technology to the desktop.

The Sun CMS will be tightly integrated with the graphics libraries in order for color management to exploit the performance advantage of acceleration hardware when it is available on a particular workstation. Integration with the graphics libraries also offers the possibility to provide accurate color to existing applications that are naive about color.

Device Characterization

In order for a CMM to reproduce accurate color on a device, it must have quantitative information about the color reproduction characteristics of the device. The collection of this information for a particular make and model of device is called a device characterization profile (DCP). Characterizing a device is a process involving science, craft and art.

CMS vendors and peripheral manufacturers are developing methods to characterize devices, and CMS vendors will provide profiles for various devices. Because different organizations specialize in different areas of technology, it is likely that CMSs from different companies will have strengths and weaknesses when used with different devices. For example, Kodak is likely to have the best characterization for their own XL7700 printer, EFI is likely to have the best characterization of the Canon CLC 500 color laser copier, and high-end prepress vendors are likely to have the best characterizations of prepress system color.

In the short term, the format of device profiles will be unique to each CMS vendor. Kodak has announced publicly that they will publish their DCP format in order to allow peripheral vendors or third parties to develop profiles for their devices, so that these devices can be exploited by ColorSense. To make it easy for naive users to get accurate color, a peripheral vendor should ship a DCP with every device. However, until agreement on DCP standards is reached, a device manufacturer will have to ship different DCPs for each CMM that he wants his device to be work with.

Device Calibration

Use of a color management system to perform transformations among device color spaces, along with characterized devices, will be sufficient for the majority of workstation applications. But devices drift, and different media have different characteristics. If really accurate color reproduction is required, it is necessary not only to have characterized devices, but also to periodically measure particular instances of devices and their media, and to make corrections. This process is known as calibration. The process of calibration appends to a DCP information specific to an instance of a device.

Device-Independent Color Image Interchange

Color management software will make it practical for users to interchange color images in device-independent color space, and thereby decouple image data from any particular device. I propose that a calibrated non-linear RGB color space can be exploited for digital image interchange in publishing to gain the advantages of device independence. In my opinion, very broad range of users — perhaps 95% of the office market and 70% of the graphics arts market — could be completely satisfied by a single well-chosen RGB space.

For a single space to achieve acceptance it must be objective, that is, it must have a tightly-defined relationship to the CIE standards for color. All of the proposed CMMs accommodate CIE spaces, so this provides a mechanism to achieve device independence.

To be practical, the chosen space must achieve very good image quality with three components of eight bits each. This means that the space must be reasonably perceptually uniform, in order to maximize the visual utility of the available codes. CIE XYZ space is not perceptually uniform, and requires about 16 bits in each component to achieve excellent image quality. Consequently, CIE XYZ is not a good choice for an image interchange space. The CIELAB and CIELUV spaces have been successfully used to convey image data in a device-independent manner using three eight-bit components. However, these spaces have the disadvantage of computational complexity.

The chosen space should be capable of good interactive performance. Users will demand reasonably fast display of newly acquired image data, without a computationally prohibitive color transform calculation. The computational complexity of spaces such as CIELUV and CIELAB prevents interactive performance.

CCIR Rec. 709 Calibrated, Non-linear RGB

These constraints suggest a calibrated, non-linear RGB space. We are fortunate to have obtained in April 1990 unanimous worldwide agreement on a calibrated nonlinear RGB space for HDTV production and program exchange: CCIR Recommendation 709. I suggest that the parameters of CCIR 709 form an excellent basis for a preferred color interchange space for digital pictures.

Note that I have been careful to avoid using the words standard color space: nothing in this approach precludes use of a different color coding if it is more appropriate to a particular application. However, in my opinion, the identification of a well-chosen default or recommended color coding has the possibility of sidestepping the creation of multiple, functionally-equivalent systems, with the attendant long-term conversion headaches.

Although they were developed for the purpose of measuring and specifying color differences, the CIELAB and CIELUV spaces have been successfully used to convey image data. Roughly the same amount of computation is required to transform any of CIELAB, CIELUV or CCIR 709 RGB to print. However, CCIR 709 was designed to be extremely efficient for driving today's CRT displays: the transform from the CCIR 709 interchange space to a CRT's RGB primaries is extremely simple, because it must be performed at real-time pixel rates in low-cost consumer equipment! CCIR 709's suitability for CRT displays does not come at the expense of applicability to alternate display technologies, however: the acceptance of CCIR 709 as the worldwide HDTV color standard ensures that it will thrive as an interchange space for new display technologies — in particular, flat panels — now under development. In addition, CCIR 709 is suitable for prepress: the ANSI IT8 committee has adopted the CCIR 709 primaries for exchange of device-independent RGB image data, although IT8 has adopted the CIE D50 white point that is ubiquitous in print work rather than the D65 standard that CCIR 709 inherited from television.

Most physical devices — scanners and CRT displays, for example — employ RGB values that are physically constrained to be non-negative. The gamut of colors that can be specified by all-positive RGB values is quite restricted. If a wide color range is to be accommodated, extensions to the traditional RGB coding range are necessary. To accommodate the wide color gamut of film or print, computer software should be designed to allow RGB values that are negative or greater than one. The coding of PhotoCD™ images uses this scheme. For example, the saturated cyan of a Salem cigarette carton cannot be reproduced on a typical CRT display: that color is outside the gamut of the CRT. However, that cyan can be represented as a set of RGB values where red is somewhat negative and blue and green are both slightly larger than unity. Upon conversion for a CRT, the out-of-range signals will be clipped, and the display will produce the most saturated cyan available. But the wide-gamut coding has the great advantage that the true color is maintained for subsequent devices, for example if the color is later to be reproduced on a dye sublimation printer, or in print, or on a photographic reproduction.

Named Color Spaces

A number of color systems are widely used to communicate absolute colors in the face of device and process instability. These systems are primarily used to specify solid colors, whether reproduced by specially-mixed inks or reproduced by four-color process printing.

The FOCOLTONE and TRUMATCH™ systems involve characterizing four-color offset printing processes. These systems have swatch books that catalog the colors that result from given percentages of C, M, Y and K. This allows a graphic designer to determine the CMYK values required to reproduce a particular color, and specify those values to a printer. However, a particular combination of values for one system's inks — say 84C 0M 44Y 21K, in Toyo inks — will produce a color similar to, but probably not identical to — the color produced when the same set of values is reproduced in a different set of inks. So each system works well individually, but exchange of colors between systems is not characterized.

The PANTONE™ Matching System identifies about a thousand unique colors by number. PANTONE provides swatch books that enable graphic artists to communicate these colors to a printer. The PANTONE system provides printers with ink formulations to produce the colors. PANTONE is most suitable for solid colors that are obtained by mixing ink rather than halftoning, but PANTONE has recently introduced a system that characterizes CMYK ink combinations, in the manner of FOCOLTONE and TRUMATCH.

The COLORCURVE™ system assigns a unique coordinate to every color. A COLORCURVE coordinate is specified by three numbers: a lightness, a red / green value, and a blue / yellow value, for example L40 G3 B2. The system is well established for specifying paint colors, and is beginning to see use in specifying inks.

All of these systems are likely to be accommodated by color management systems as plug-in color spaces.

Application Program Interface

In the past, color transformation capability has been provided directly by several specific applications. However, the developers of nearly all application programs would prefer color to be handled at a system software level. The applications that have offered their own color transform capabilities have done so mainly in order to offer interfaces to high-end color printers — and to prepress systems — in advance of color management being a built-in feature of the workstation environment.

Most applications are transparent to the color interpretation applied to their image data. A program that reads, writes and manipulates RGB data will obviously work with any interpretation of the data. Interposing the CMS between the application and the graphics hardware — scanner, monitor and printer — will allow application to inherit accurate color capability from the CMS.

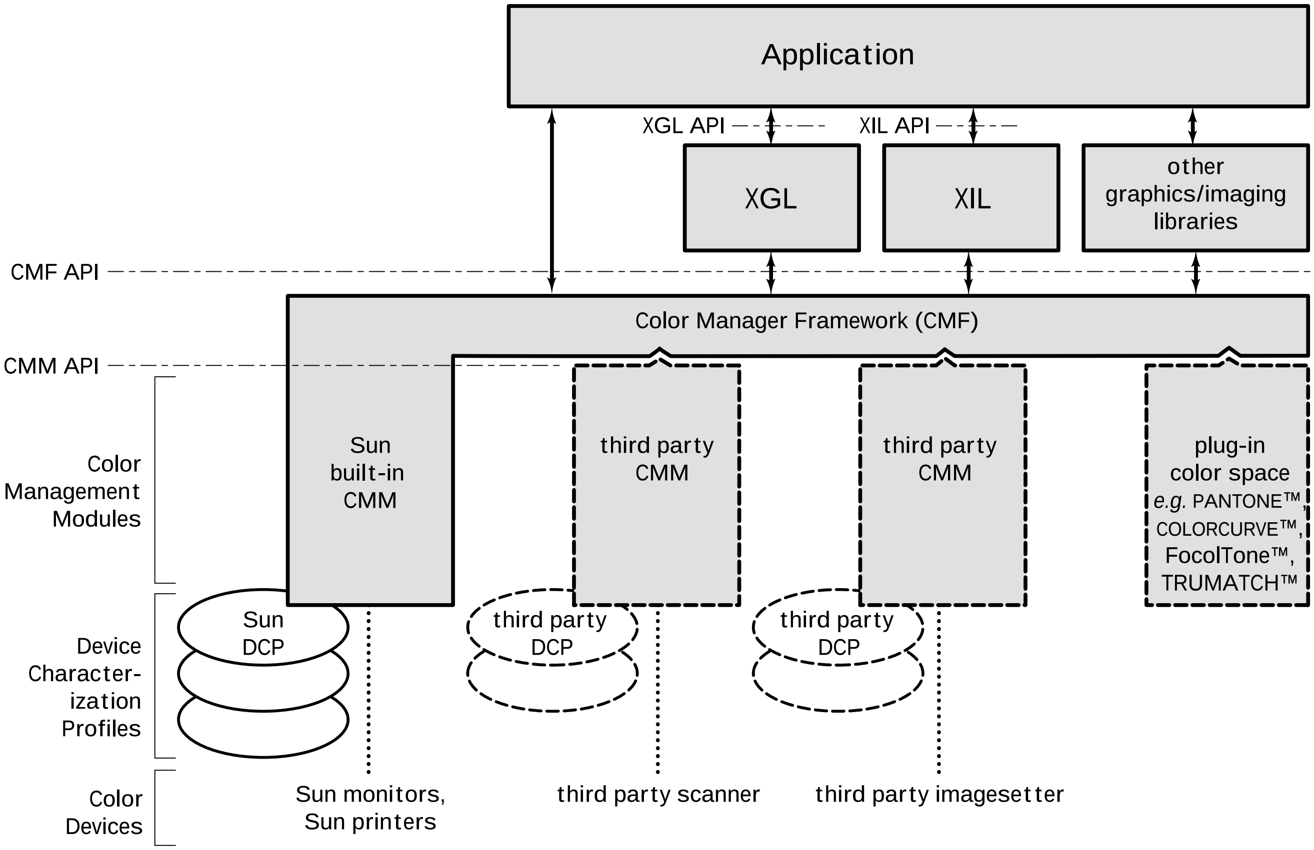

Figure 1 Color Management System (CMS) capabilities are accessed by an application program, illustrated at the top of this block diagram. Underneath the application is a set of graphics libraries, each presenting an application program interface (API). Underneath the graphics libraries is the color management framework (CMF), which serves as a dispatcher for color capabilities that are presented to the application and to the graphics libraries by the CMF. The mathematical transformations of color are performed by color management modules (CMMs) that plug into the CMF through a private CMM API. To perform a transformation, a CMM needs access to a device characterization profile (DCP) specific to an input or output device: a scanner, monitor, printer, imagesetter, etc.