Standards Update

Orchestrating Color

Expectations & Understanding Vary Widely

by David McDowell, NPES/Eastman Kodak

In today's world of color management we find ourselves in an interesting situation. We have too many options and capabilities, there is no dominant solution that is providing a common direction, and very few users understand the basic principles of color management.

As a result, user expectations and understanding vary widely and often conflict both with one another and with the capabilities of existing applications.

It is easy to see why this situation exists. The standards provide the file formats to enable the exchange of data—but don't regulate exactly what data should be exchanged.

The color management vendors provide the tools to manage color but leave it up to the user to decide what to manage and when to manage it. The page composition programs also leave it up to the user to decide what to do. The picture editing programs are even more flexible and allow the user complete freedom to edit in any color space desired.

The closest thing we have to any applied industry systems architects are some of the industry consultants, and each has his or her own agenda.

Color management gives us tools and capabilities that allow us to be more efficient and to do things that were impossible only a few years ago. Let's use what we have and be patient about today's limitations because they will be solved tomorrow.

Recent History

To understand the impact of all of this, let's first look back at recent history and see how color was managed a few years ago.

As we look back to the days of CEPS (Color Electronic Prepress Systems), 1982-1992, the vendors provided the system integration and workflow solution. CMYK was the data format used for both editing and exchange and there were few variations from system to system. Images were scanned directly to CMYK and most global corrections were done during scanning. These included cast removal, tone reproduction adjustments, gamut compression, overall color correction, image sharpening, etc. Differences in color correction techniques between different brands of scanners was often exploited by trade shops to help optimize images for different applications. Workstations were used to accomplish local color editing and correction as well as any necessary fine tuning of the global corrections. Text and line art were also handled, as scanned data or as composed raster data. Spatial editing (cropping, artifact removal, image enhancement, drop shadows, etc.) and image assembly were all done using raster files.

The desktop systems that followed them (1987-present) were also predominantly CMYK workflows—at least at the point of data exchange. The desktop systems in many ways mimicked the workflow of the CEPS systems, using newer, faster computers. However, because of their increased computing capability, they did begin to do some of the global corrections in the com-

Standards articles published by the IPA Standards Committee. Editors: David McDowell, NPES/Eastman Kodak. and Larry Warter, Fuji Photo Film.

Orchestrating Color *is challenging because we have so many options and capabilities and there is no dominant solution providing a common direction, plus very few users understand the basic principles of color management. As a result, user expectations and understanding vary widely and often conflict with each other and with the capabilities of existing applications.*

puter instead of the scanner. Input profiles were used along with packages like Photoshop to convert raw scanner data to intermediate RGB or CMYK. This period also saw the introduction of PostScript in the graphic arts and the ability to treat text and line art as vector data (at least during some parts of the workflow). There was also some experimentation with RGB workflows—mainly for image data.

During this period, TIFF/IT (Tag Image File Format/Image Technology) was the dominant file format for data exchange and any indication of the intended printing condition was carried outside of the content files.

In the CEPS and desktop workflows, during much of this time, image color was corrected and proofed as individual images prior to page assembly. CGATS TR001 arrived on the scene in the mid 1990s as the digital representation of SWOP (Specifications for Web Offset Publications) proofing but was mainly used to set up separation aims and/or color proofing devices. Image data were edited with the intended printing condition in mind and the printer expected to receive CMYK data and a hard copy proof.

Although the IT8 scanner targets were introduced in the early 1990s, they were initially used along with a form of a CMYK destination profile to help convert data into CMYK in a variety of input tools—many of them Photoshop plugins or similar packages. These, in truth, represented some of the early applications of color management, albeit in rather closed environments. Many of the scanner characterization tools really prepared a form of input profile using the same principles as today's color management architecture.

Initial Graphic Arts Use of Color Management

Two of the earliest applications of output color management were the continuous tone color proofing devices and gravure cylinder engraving. These both compared the color that was expected from a CMYK offset file—for example CGATS TR001 for SWOP—and created the data conversions necessary to match gravure to offset or proof to print.

At the same time, applications were being developed that allowed end-to-end color management based on the ICC (International Color Consortium) architecture model. However, without clearly defined standards for both a workflow architecture and file formats to convey the necessary information, they were largely constrained to in-house applications. These focused on delivering CMYK data for exchange more efficiently than traditional systems.

Within much of the world, CMYK data are still the basic exchange format for graphic arts data. However, color management is more and more being used to create that CMYK data. Even when color management is not used in creation of CMYK data, it is being used to identify printing conditions for which the CMYK data were intended.

The Color Management Model

One of the problems has been a misunderstanding of the basic color management model as it was initially conceived. It started as a device independent model to decouple sender and receiver of office documents and pictorial information. The basic premise was that the sender should include the most information possible about the content but it was then up to the receiver to do the best possible job of reproduction with the resources available at the receiving site. All the sender could indicate was a desire for the receiver to emphasize perceptual or colorimetric reproduction or to preserve saturation. Because the ability to provide perceptual rendering was included, many in the color management community felt this was all that was required to ensure good quality results. The common understanding of perceptual rendering was to do the best job possible of making the reproduction look like the original within the capabilities of the output device.

This was widely touted as the way to answer all of the color needs of the image reproduction business and lavish promises were made about how well it would work. It does meet many of the needs of the information and office systems world, including many web publishing applications, and it did decouple the sender and receiver. It also made it impossible for the sender to either predict or have any control over the final reproduction.

This model does not fit many of our graphic arts applications. In graphic arts, the sender (read prepress provider) is usually expected to get customer approval of a proof of the expected (desired) reproduction. The sender is expected to understand and/or define the specific conditions under which an image will be reproduced, regardless of the output device chosen by the receiver (read publisher or printer). This disconnect between the basic ICC architecture and the needs of the printing and publishing industry has led to many of the misunderstandings that exist today. The ICC is now addressing the needs of our industry but this has required major changes in the ICC perspective.

Another issue that leads to a great deal of confusion is that the ICC color management model does not address where edits and corrections take place. In a fully implemented color-managed workflow, input data (typically scanner data for images) are expected to flow through input and output profiles, under control of the CMM (Color Management Modular), and be converted into the intended output data (CMYK for graphic arts) in a single processing step. Prior to processing, input data are stored with the appropriate profile until the ultimate destination is determined.

In this classic color management model, the opportunities to edit the data are to: go back and edit the input scanner data; to adjust the scanner setup; to edit either the source or destination profiles used; to add effects profiles; or to edit the final CMYK data created. Other than editing the final CMYK data, few applications provide the ability to accomplish these edits.

As a result there is a growing tendency to use an intermediate RGB space in which input data can be stopped and held for editing. This, in essence, creates a two-step color management process. Color manage to an intermediate color space where data are edited and assembled, and then color manage from this intermediate space to the final output. If this intermediate space is an 8-bit per color encoding, there is a high potential for the introduction of artifacts.

One other aspect of the color management model we need to briefly review is the role of the source and destination profiles. In most workflows the role of the input profile is relatively simple. That is convert device (scanner) code values into PCS (Profile Connection Space) reflection color data and compensate for differences in viewing conditions (transparency to reflection, etc.). The destination profile, on the other hand, is required to do a lot of work—particularly the perceptual transforms going to CMYK data. Virtually all of the work that used to be done in the scanner is done in the output profile—gamut compression, tone scale adjustments, color separation, black printer generation, white point adjustment, etc. Editing could also be done using effects profiles or device link profiles but few applications enable their creation and none provide the tools to link them to a specific image.

Unfortunately, even in PDF/X-3 (the most color management friendly of all of the imaging file formats) only one profile can be associated with each image and that is the source. The PDF/X file format itself does allow a destination profile to be carried, but that profile must be used with all of the color-managed data in that file. Therefore, if there is more than one color-managed object in a file, they all must use the same destination profile.

How Is Color Management Being Used Today?

CMYK Data Exchange

Within much of the world, CMYK data are still the basic exchange format for graphic arts data. However, color management is more and more being used to create that CMYK data. Even when color management is not used in the creation of CMYK data, it is being used to identify the printing conditions for which the CMYK data were intended. The PDF/X-1a file format (defined in ISO 15930-1) requires pointers to standard characterization data, to be included as part of the file. The preferred registry that is identified in the PDF/X standards is the ICC registry at **www.color.org**. Where the expected printing does not match a registered characterized printing condition a destination profile must be included.

This has placed much more emphasis on the ICC characterization registry and characterized printing condition data that are identified in that registry. For the first time in the history of our industry, we have a single location to point to where established sets of data that relate CMYK input values to printed color are identified.

Also, for the first time, our data exchange formats require such information. Admittedly, there has been some naming confusion in the ICC registry, but TC130 and the ICC are working together to get that straightened out.

Because a color management system that delivers CMYK data typically does not interact with other color management systems, many of the current color management issues, while frustrating, do not have an impact on the final CMYK data exchange. The sender still can look at a CMYK proof before committing the data to a printer or publisher.

Exchange of Color-Managed Data

There is increasing interest in exchanging three-component data, more so in Europe and newspaper applications. Most people casually refer to this as exchanging RGB data, forgetting that there are many RGBs from which to choose. Here the relationship between color management and data exchange is closely coupled.

In September 2002, the first graphic arts data exchange standard that fully enabled the exchange of color-managed data came into existence. PDF/X-3 (ISO 15930-3:2002) represents a major step forward and allows the exchange of fully defined three-component data for graphic arts applications. It requires the use of ICC destination profiles to identify the intended output condition and to define the data conversion between the PCS and the input code values of the intended printing device. It also makes provision for source profiles to be used to define that specific three-component data (RGB) being exchanged. However, it does not say what three-component data should be used nor does it provide profiles, etc. These are all user choices.

The same application areas that are encouraging the exchange of three-component color-managed data also seem to be increasingly accepting soft proofing on the color monitor.

Let's look at some of the issues (and potential pitfalls if not handled properly) involved in exchanging three-component color-managed data based on soft proofing.

Display-The display profile (yes, in addition to being well-controlled and calibrated the display device must have a profile to convert data from the PCS to the display input values) must compress or clip data so that it will fit the gamut of the display device. The typical monitor is approximately sRGB. Typically an intermediate work space will be a large gamut RGB (always larger than sRGB). Also, although most people do not realize it, in some parts of color space typical CMYK printing has a larger gamut than sRGB. On top of all of that, appearance modeling must also be used to make a relatively dim self-luminous display "look" like a reflection print under high illumination.

The old term WYSIMOLWYG, *what you see is more or less what you get*, really applies to soft proofing. Yes you can learn to compensate and estimate but it takes experience.

Profile Interchangeability-Source and destination profiles are based on proprietary technology. Profiles from one vendor will NOT produce the same results as those from another vendor, nor should they be expected to. Some of those differences are what allow vendors to differentiate themselves. Different perceptual CMYK destination profiles, even from the same vendor, may handle tone reproduction, gamut compression, and black generation significantly different.

That is why PDF/X-3 says that the profile included should be used to render the data to CMYK. Too many people believe that if they decide to change output devices they can simply substitute a new profile and get similar results. It doesn't work that way.

Even with colorimetric profiles, different colorimetric profiles should produce colorimetric values that are close to each other, but they all handle colors near the gamut limit differently. In addition, going from PCS to CMYK data, each vendor has unique color separation and black generation algorithms—color should be close but the components will be different.

Image Assembly-The issue of the assembly of multiple files using three-component color-managed data has not been cleanly solved by the standards community or by the application vendors. We must associate a source profile or color space definition with each object. However, we cannot associate any other profile with individual objects. There is one destination profile that applies to the whole file. If we want to treat images differently within the same file, e.g., high key vs. low key tone reproduction in a destination profile, we cannot do that in a three-component color-managed workflow. Further, if two files are prepared for the same characterized printing condition, but use different output profiles (or profiles from different vendors) they cannot be combined without additional processing.

The caution in the PDF/X-3 application notes says, "If device-independent color data is used in PDF/X-3 files, the profile included in the Output Intent of each file must be compared to those in all other files to be assembled together. Where all profiles are identical, the files may be assembled directly, retaining device independent colors. If different profiles are used, then colors must be transformed to the output device color space prior to assembly to ensure that the correct gamut and tone compression is performed for each entity." Currently there are no other obvious solutions!

Planning Ahead

IPA Technical Seminar
May 13 - 15, 2003
Holiday Inn Mart Plaza
Chicago, Illinois

Graph Expo & Converting Expo
September 28 - October 1, 2003
McCormick Place South
Chicago, Illinois

IPA Management Conference
October 12 - 13, 2003
Pointe South Mountain Resort
Scottsdale, Arizona

Black Channel Preservation-The classic color management model says that to convert CMYK data from one device to another (where the gamuts are the same or close to each other), combining a colorimetric device to PCS transform for the first device with the colorimetric PCS to device transform for the second device, will yield the correct colorimetric results. And it will, except that the color separation scheme and black printer will be what was included in the profile for the second device and may not bear any relationship to the initial CMYK. If this is for a CT proofing device, it is probably more than acceptable. If the black-to-color relationship is important, then some other transform is required—a number of applications have the ability to create black preserving device link transforms.

This is the classic problem that is faced by proofing systems and those systems that want to optimize CMYK data for a specific output device. Here the gamuts are correctly maintained by process control of solid ink density, but differences in tone value increase, trapping, etc., mean that different CMYK input is required for within gamut colors. Using the gravure process to match offset SWOP data is a perfect example of this situation.

Repurposing-Repurposing, not to be confused with retargeting, is sending output to a device with a different gamut than the gamut it was initially prepared for, for example, CMYK publication data to a web display. Retargeting is sending data to a device with the same gamut but different encoding.

In repurposing, the first decision that must be made is whether the appearance in the initial output mode (e.g., CMYK publication) should be preserved. If so, the output data must be colorimetrically converted back to PCS and then either a colorimetric or perceptual output profile used to convert to the new destination, depending on the relative size of the color gamuts of the initial and new destinations.

If the appearance in the initial output is not significant, then a new destination profile can be substituted but the image should probably be reproofed for the new output condition to be sure the intent of the designer is preserved in the new output color space. Too often users believe that simply substituting a new destination profile without additional proofing or checking will produce comparable results on a different device. It is simply not that easy.

Where Do We Go from Here?

Color management does work and three-component data exchange can produce reliable results. There are limitations and cautions that must be observed.

Some have suggested that if all color management systems did the same thing—used the same gamut compression, tone reproduction, etc.—many of our problems would go away. Others have suggested that images should be gamut compressed and have tone reproduction corrections applied before being placed in a color-managed workflow so that only colorimetric transforms would be required. These would allow much more commonality between applications but at the expense of flexibility and repurposing. This approach would make color management a lowest common denominator system with few incentives for vendors to participate.

We have come a long way; there is much yet to be done! Today color management gives us tools and capabilities that allow us to be more efficient and to do things that were impossible only a few years ago. Let's use what we have and be patient about today's limitations because they will be solved tomorrow.