

Limitations in Communicating Color Appearance with the ICC Profile Format

Gary Dispoto and Michael Stokes
Hewlett-Packard Company, Palo Alto, California

Abstract

The ICC format has been widely adopted as an industry standard for communicating color and a context has been presented for implementing and interpreting this standard. There are three major implications of the way the ICC profile format has been publicly presented for use in communicating color; user interface and application development complexity, possible quality limitations due to undefined transforms, and the possible quality degradation of business graphics. Some specific implications for peripheral vendors will also be discussed.

Review of Published Material

In attempting to implement and interpret ICC profiles, there is a limited amount of available documentation. This includes the official ICC specification¹ and one paper by Kodak² (one of the ICC founding members) on implementing the ICC profile format. A brief summary of key points relevant to this paper from the above references are provided below.

Section 2.6 (Profile Connection Spaces) of the ICC specification summarizes the justification for and definition of the ICC Profile Connection Space (PCS). The summary emphasizes the need for a “well-defined” and “unambiguous” connection protocol between ICC profiles. A well-defined PCS allows device profiles to be connected together transparently providing a transformation of colors from one device to another. Classical CIE colorimetry describes color matching under similar viewing conditions. Extensions to classical CIE colorimetry are necessary to compensate for different viewing conditions or media. The PCS is defined as “the CIE colorimetry which will produce the desired color appearance if rendered on a reference imaging media and viewed in a reference viewing environment.” The CIE colorimetry is referenced to the graphic arts standard ANSI CGATS.5-1993, “Graphic Technology—Spectral measurement and colorimetric computation for graphic arts images.” The reference imaging media is defined to be an “ideal reflection print.” The reference viewing environment is referenced to the graphic arts standard ANSI PH2.30-1989. It is the responsibility of the profile builder to “adapt the measured colorimetry to that *appropriate* for the” PCS. Variables included in this adaptation are; absolute media white point, illuminant white point and illuminance level, media black point, white point and luminance level of the viewing surround, and flare.

Section 3.2.2 (RGB Display Profiles) of the ICC specification details the required tags necessary to include in an RGB display profile in an apparent contradictory manner.

The red, green and blue “colorant” tags describe the relative XYZ values for the display phosphors. This matrix of XYZ values converts the “linearized values into XYZ values for the CIEXYZ encoding of the profile connection space.”

Appendix D (Profile Connection Space Explanation) of the ICC Specification is “intended to clarify certain issues of interpretation in the” ICC profile format. This appendix re-emphasizes the previous definitions of the PCS given in section 2.6 of the ICC specification. The goal of this appendix is to provide an unambiguous connection protocol for device profiles. This “*device-independent specification* is best given in a color space based on human visual experience. Thus, a device profile provides a means of translating (or transforming) color image data from device coordinates into a visual color space or vice versa.” The PCS is composed of “a coordinate system for the color space and an interpretation of the data represented in that coordinate system.” This interpretation not only defines the CIE colorimetry, but the media and viewing conditions. These conditions include; absolute media white point, illuminant white point and luminance level, media black point, white point and luminance level of the viewing surround, and flare. A media “personality” is defined to include the properties of dynamic range and color gamut along with the methods of mapping out of range and out of gamut colors into the media dynamic range and color gamut. The “desired” part of the PCS definition “implies the expression of artistic intent.” “It is essential to the success of color-management systems that a broad range of options be kept open. The interpretation of the PCS merely defines the particular default behavior that will be facilitated by the system without explicit intervention by the application or user.” The appendix goes on to describe a series of output and input device scenarios and which interpretation variables must be accounted for in each situation. “If the creators of device profiles universally apply these corrections to their colorimetric data, the PCS will have a universal, unambiguous interpretation, and image rendered ‘colorimetrically’ will evoke (as nearly as possible) the sample appearance regardless of the medium or viewing environment of the reproduction.”

At this conference last year, Kodak presented a poster paper “Communicating Color Appearance with the ICC Profile Format.” The goal of this poster paper was an unambiguous definition of a PCS that connects device profiles. The definition of the PCS is described to be a virtual space that quantifies how an image appears in order to communicate it and ultimately allow its reproduction. The concept of the colorimetric context is defined to be the

extensions to colorimetry that allow the preservation of color appearance “across different media and different viewing environments.” A colorimetric context includes; absolute media white point, illuminant white point and luminance level, media black point, white point and luminance level of the viewing surround, and flare. These variables are described in some detail. The colorimetric context in this poster paper appears identical to the *PCS interpretation* of Appendix D of the ICC specification. [Note: The term *colorimetric context* will be used below to describe this concept.]

The distinction between a *narrow PCS* and a *broad PCS* is based on whether the colorimetric context parameters are fixed or variable. “If colors in the PCS are to be interpreted as rendered on a medium with certain fixed properties then the media related aspects are narrowly defined.” Thus in a narrow PCS, the colorimetric context parameters are fixed to implicitly agreed upon values. This forces all profile properties to be statically defined and fixed when the profile is created. A *broad interpretation* is “an equally valid alternative to associating fixed media properties with the PCS is to explicitly communicate information about the media through data fields in the profiles. This leads to a broadly defined PCS—one without a strongly associated reference medium.” This interpretation allows profile properties to be dynamically changed to reflect variances in the device or viewing environment at any time.

“The issue of whether the PCS is defined narrowly (with a fixed context) or broadly (with a variable context) is fundamental to the design of a color management system. The resulting choices lead to quite different designs for the profiles and CMMs” (color management modules). The advantage of a narrow PCS is “that obtaining good cross-media mapping relies less strongly on knowledge of the input and output media. This approach facilitates applications which place more emphasis on “pleasing” reproduction than strictly colorimetric reproduction.” Another advantage of this approach is that the CMM can accommodate profiles from differing media and environments without any extra functionality (essentially the CMM can simply be an interpolation engine). Finally, a narrow PCS increases interoperability of a given profile between different operating systems. “Since the CMMs will have to run even on low-end computers, it is not possible to obtain both acceptable quality and acceptable real-time CMM performance.” The advantage of a broad PCS is that it is “easier to render to properties (or “personality”) of one medium on a medium with different properties (although it is still possible to obtain “pleasing” renderings with some extra work.” A broad PCS also facilitates interoperability of profiles from different vendors on a given vendor’s CMM. “The most important point to make is this: ICC profiles are conceived of as “ready to use” transforms that can be unambiguously plugged together by a CMM that does no modification of colors in the PCS.”

Expectations About the ICC Profile Format

One of the expectations about the ICC profile specification is to be able to build a profile from it. Like a cooking recipe, this implies a step by step approach can be derived, all of

the required ingredients are described as well as how to combine these required ingredients together. Also, like a cooking recipe, “value add” might be interpreted as biasing the recipe to a particular taste or audience. This is not the case with the current ICC profile format specification.

Interoperability is defined to be how well objects work together and how much work the object vendor must perform in order to accomplish this. The burden of interoperability should be placed on the fewest vendors (who hopefully have the most resources). This would be the operating system (OS) and color management framework (CMF) vendors. The next burdens should be placed on the CMM vendors. The profile vendors should have fewer burdens than the CMM vendors and the applications vendors should have fewer burdens of interoperability still (with respect to color management). Finally, the users should be the least burdened by interoperability issues. This clearly places the operating system and color management framework vendors in a position of responsibility and leadership in the field of color management. It currently appears that the ICC has implicitly placed greater importance on interoperability of CMMs than profiles and placed the greatest burden on profile vendors.

ICC Colorimetric Context

It appears that there are two distinct philosophies competing within the ICC specification. One is a narrow PCS with a fixed colorimetric context and the other is a broad PCS with a default (but variable) colorimetric context. Both approaches use the same clearly defined set of colorimetric context parameters with well-defined default values. The difference is that a narrow PCS only allows these default values, while a broad PCS allows optional tags to override these default values.

Since no explicit default values for the PCS colorimetric context have been clearly stated in the ICC specification, we propose the following values;

1. viewing illuminance: 500 lux
2. viewing white point: CIE D50
3. media white point: that of a perfectly reflecting diffuser
4. media black point: 0 lux or 0 Reflectance
5. flare: 1 percent
6. surround: 18.4 percent of the media white point
7. media description: reflection print (i.e. RLAB, Hunt viewing media)

These parameters are consistent with those of the Hunt, RLAB, and Nayatani color appearance models. It is imperative that the default values of these parameters be numerical values and not loosely defined ranges of values. While an ideal print has no flare, the one percent flare value was arbitrarily chosen to be consistent with current print standards. Additional media descriptions include; 35mm slides, cut-sheet transparencies, monitors, and self-luminous colors.

The ICC concept of a reference imaging media in the colorimetric context as “an ideal print” is confusing. The term “ideal” in this case implies unlimited dynamic range (i.e., black point is 0.0 luminance and white point is a per-

fect reflecting diffuser), an unlimited color gamut and no flare. This is much closer to most color appearance models than any real media and would seem void of a media personality. Aside from the additional media description, it is unclear what this adds to the colorimetric context.

Another concern is that the ANSI CGATS.5-1993 measurement standard requires black backing for all measurements. This is contrary to current industry practice and problematic in practice. Measurements on black backing are NOT representative of typical use and difficult to transform to a representation of typical use.

Related Colorimetric Context Tags

There currently do not appear to be public tags to specify either flare or a media description that are consistent with most modern day color appearance models.

The fact that there exist optional tags for context parameters which are fixed is confusing. This would imply that the tags are redundant or only informational, but the ICC specification does not provide any guidance on how to use these tags appropriately.

The RGB display profile tag descriptions in section 3.2.2 of the ICC specification imply conflicting standards. The colorant tag descriptions state that they are XYZ values of the phosphors, but on the same page it is stated that they are the XYZ values of the CIEXYZ encoding of the PCS. This implies the colorant tags are possibly used for storage of both phosphor data and colorimetric context data. The individual tag descriptions again refer only to the phosphor or colorant and not to the PCS. If the PCS context is not part of the XYZ tags, the specification is silent on where to incorporate the colorimetric context parameters in the CRT model. The required ICC CRT model is a simple 3×3 matrix multiplication followed by three one-dimensional lookup tables. This is consistent with previous CRT modeling research for transforming from device RGB values into CIEXYZ values. In order to transform into the fixed D50 PCS, one must also use these parameters to account for white point, viewing conditions and flare. One possibility is to use a simple Von Kries white point adaptation folded into the 3×3 matrix parameters and fold the flare and surround effects into the three one-dimensional lookup tables. While this provides a possible solution, it makes extracting CRT device characterization data, such as phosphors or gamma values, nearly impossible without detailed prior knowledge of how the profile was created. This seems in conflict with the goal of having an open device characterization profile format and also the naming and descriptions used within the ICC specification itself.

Process

One shortcoming of the ICC specification is the incomplete description of a processing pipeline. While some parts of this pipeline are described in extreme detail, such as sections 6.4 and 6.5, other parts such as sections 3.1.2 and 3.2.2 leave significant portions of the process as *an exercise for the reader*. It would be helpful if the ICC provided a baseline algorithmic process description for the entire pipeline instead of only portions of the pipeline. While it is understandable that this is an opportunity for companies to

profit by their expertise in this area, the industry as a whole would benefit tremendously from a complete baseline process description which could be enhanced as vendors see fit. This would involve completing what was started in the ICC specification. One should be able to read the ICC specification and build a complete color transformation between two devices from mathematical equations provided by the ICC specification. This is consistent with the approaches of OpenGL, X, the Common Desktop Environment and other cross-platform initiatives and seems imperative if the ICC intends for this specification to become an ISO standard.

The ICC documentation doesn't specify any particular method for transforming into and out of the ICC PCS. In fact, Appendix D, specifically states that the transformation can be based on any number of color appearance methods and "the PCS will have a universal, unambiguous interpretation, and images rendered "colorimetrically" will evoke (as nearly as possible) the same appearance, regardless of the medium and viewing environment of the reproduction." The implication is that combining profiles with different methods for transforming into and out of the ICC PCS will seamlessly work together. This is in direct conflict with many published results showing the incompatibilities between different color appearance models.^{3,4,5} The Kodak paper states that "if profile vendors take radically different approaches, users will have to use profiles from a single vendor for best results." This would limit, in direct contradiction with the former statements, the use of sophisticated color appearance models such as Hunt, RLAB or Nayatani. It also brings into question the "openness" of a specification which mandates conformance with unspecified proprietary solutions.

Limitations of the Current Narrow PCS Interpretation

The current ICC specification with its fixed PCS requires profile builders to convert between the device context and the PCS context. This can be accomplished by using an available color appearance model, such as Hunt, Nayatani, RLAB, LLAB, or others. This transformation involves converting the device colors in the device context into a human visual system (HVS) context-free representation and then into the PCS context. Therefore transforming a color between two devices involves transformations from a device context to a HVS context-free representation to a PCS context then onto a HVS context-free representation and finally to the target device context. While, in theory, these transformations can be concatenated together without an impact on quality, in practice, this is often not the case and perceptible quantization and other image degradations can occur. In summary, the current narrow PCS requires four implicit context changes to convert colors between two devices, while the broad PCS requires only two context changes in the same situation.

Another limitation with the narrow PCS is the apparent inability to edit or calibrate monitor profiles. Monitor calibration involves measuring and modifying the phosphor parameters and resetting the white point by a different normalization matrix. The narrow PCS hides the monitor phosphors by confounding colorimetric context information into the 3×3 matrix. Therefore, unless one has prior

knowledge of the profile creator's (and any editors') adaptation methods, it is impossible to reliably extract the monitor phosphor information from the profile and calibrate the monitor. This is a hindrance to the current state of color management utilities today.

The narrow PCS has significant limitations with respect to user interface issues.

Having a profile with static properties implies that a separate profile must be created for each combination of device setups and viewing environments, such as black generation, paper types, and lighting conditions. If a user has access to multiple devices (i.e. on a network) with multiple setups, the number of individual profiles can quickly grow to unwieldy proportions. A conservative example of this is a user with an ink jet, a dye sublimation and a color laser printer, each with three paper types, two black generation settings, two halftoning methods and two viewing conditions. This user has 72 profiles just for these three printers. Obviously, the addition of other types of devices, such as scanners, monitors, cameras or other printers, significantly worsen the situation.

This creates a very serious user interface problem. Currently, this is not addressed in any of the ICC profile documentation. A guideline for naming profiles would be helpful. If a broad interpretation is used, this situation is lessened due to the fact that each of the properties can be dynamically defined. In the above example, the user could have either nine or eighteen profiles (depending on whether the profile data is sufficient for different halftoning methods, given the base color transformations). While this is a significant improvement from the previous example, it is clearly not desirable for the user to have to pick from a myriad of profiles on the screen. This leads to the need for intelligent profile filters in either the CMF or in each and every application. If this is incorporated into the CMF, the applications must submit a unique set of parameters in order to get a single profile back or the applications must choose (either explicitly or implicitly) between multiple profiles meeting these criteria. If unique combinations are not possible and the user must choose the correct subset of profiles, then a convention for describing profiles is desirable. The large number of profiles, places an extra burden on the application and device driver vendors.

Depending on the size of each profile, the difference in numbers of profiles can produce a significant difference in available disk space. In the examples above, with a typical 45 kilobyte printer profile, this leads to a difference between 400 or 800 kilobytes for the broad interpretation versus 3.2 megabytes for the narrow interpretation.

Perceived Limitations of a Broad PCS Interpretation

One perceived limitation to this approach is that enhancing the baseline pipeline would by necessity create different quality results than the default pipeline. This is no different than the current situation despite all claims to the contrary. Because different vendors use different transformations to convert into and out of the ICC PCS from the device context, combining profiles from different vendors creates unpredictable results. This is actually an advantage of the

broad PCS in that profiles from different vendors use a single CMM when transforming colors so that the transformations will be consistent with each other.

There is concern over maintaining acceptable CMM performance if any intelligence at all is integrated into the CMM. Since at least two device profiles are involved in color transformations, the concatenation pre-processing can be enhanced to compensate for different colorimetric contexts. While this will significantly slow the pre-processing, the vast majority of the color transformation performance is at in the processing and not in the pre-processing stage. The point being that for a 4 bit run-time lookup table, one would only have to pre-process about 17K color values, which, even for complex models, is not a significant amount.

This does not, in any way, relieve the device vendors from creating superior profiles for their devices. It seems obvious that the device vendors would know the most about how their device works and therefore be in the best position to build a profile for it.

One consequence of attempting to evolve the ICC from a narrow to a broad interpretation is that CMM software would need to be revised in order to handle dynamic profile properties. Fixed properties profiles would continue to be valid. Therefore, the burden of implementing a broad interpretation would rest on the few CMM vendors and not the profile vendors.

Business Graphics

In attempting to manage business graphics (or solid color text in images), it is desired that pure colors on the monitor to map to pure colors on the print. At first glance there seems to be some ambiguity in the ICC specification. On the one hand, the saturation intent seems geared toward this very thing and on the other hand the concept of the PCS seems to make any reference to source device primaries incommunicable. There seem to be four possible solutions; 1) turn the CMF off and use device dependent transformations, 2) reverse engineer what the profile vendor did and then compute new original primaries, 3) incorporate a separate device dependent pipeline within the ICC architecture, and 4) map large areas of the three-dimensional component of the printer profile to pure colors. All of these solutions have significant weaknesses, some of which are unacceptable for most consumers. The ICC should clarify the use of the saturation intent and how business graphics should be handled using ICC profiles.

Exposing Profiles

It is not clear from the ICC specification whether ICC printer profiles must be "exposed" as publicly accessible or whether they can be embedded within the printer driver software and private to the printer driver. Advantages of exposing the profiles include enabling printer previews independent from the printer driver and enabling printer profile calibration for changes in ink, media or environmental characteristics. A major advantage of embedding the printer profile is performance.

Proposed Changes and ICC Scope

There is also concern that all modifications to the ICC specification be fully and completely backward compatible. While broad PCS profile would be binary compatible with a narrow PCS profile, the OS, CMF and CMM vendor would have to modify their source code and provide upgrades to process broad PCS profiles. The OS vendors of the ICC have repeatedly stated that it is their desire to provide an open, public solution to resolve the most common color management issues.

The ICC needs to examine the scope of the ICC profile format specification. If it is strictly a profile format, then the narrow PCS and broad PCS interpretation only involve how the data is interpreted and not binary changes to the profile format itself. The interpretation of the PCS widens the scope of the ICC profile format to impose standards on the CMM and color management community as a whole. This issue of scope needs to be clarified, or else the industry will end up with systems designed to conform to data objects and not as fully engineered software systems. If this is the case, then ICC is not just setting a profile format standard, but is setting system level standards for how colors are managed within and across computer platforms.

Conclusion

The ICC profile format provides a strong solution for communicating color information. We believe that this stan-

dard could be made stronger yet by changing from a narrow to broad interpretation, which would resolve many of the limitations described above. Although only the fixed vs. variable colorimetric context issue has been discussed in detail, other significant issues which limit color communication using the ICC profile, but are beyond the scope of this paper include; 1) the disparity between the graphic arts and photography perspectives on color management, 2) incomplete measurement method descriptions, 3) vaguely defined rendering goals and intents, and 4) the scope of the ICC specification.

References

1. InterColor Consortium, "InterColor Profile Format", version 3.0, June 1994.
 2. S. Gregory, R. Poe, D. Walker, "Communicating Color Appearance with the ICC Profile Format", *Proc. of IS&T/SID 2nd Color Imaging Conference*, 1994.
 3. M. D. Fairchild, "Visual Evaluation and Evolution of the RLAB Color Space", *Proc. of IS&T/SID 2nd Color Imaging Conference*, 1994.
 4. M. R. Luo, et. al., "Quantifying Colour Appearance: Part IV. Transmissive Media", *Color Research and Applications*, Vol. **18**, No. 3, June 1993.
 5. M. Mahy, et. al., "Evaluation of Uniform Color Spaces Developed after the Adoption of CIELAB and CIELUV", *Color Research and Applications*, Vol. **19**, No. 2, April 1994.
- ☆ This paper was previously published in *IS&T's 3rd Color Imaging Conference Proc.*, p. 155 (1995).