

Communicating Color Appearance with the ICC Profile Format

*Scott Gregory, Robert Poe and Douglas Walker
Eastman Kodak Co., Billerica, Massachusetts*

Abstract

The advantage offered by the ICC system over proprietary color management systems is its potential interoperability. This not only requires agreement on a file format but also on the way this format communicates appearance. The successful use of the format requires both cross-media and cross-environment transformations. To achieve interoperability, the ICC format places these transformations in the profile rather than in the color management engine. This paper does not attempt to summarize the ICC specification. Instead, it establishes a broader context for understanding the way the design of the ICC communicates color.

Introduction

A consortium formed by a group of major companies active in the digital imaging business has recently published an open file format which is likely to have an impact on the way that color is communicated. The format, called the International Color Consortium (ICC) Profile Format,¹ allows a standard representation of *device profiles* as used by color management systems for describing color transformations. Although this format is not that different from private ones used in a number of existing color management systems, the desire of these major companies to be able to communicate color in a universally accepted manner poses new challenges. Central to these challenges is the task of representing the appearance of color images as viewed on many different media and under very different viewing environments. This paper surveys the requirements for communicating color and discusses the practicalities of implementation in a cross-platform environment.

The Goal of Interoperability

The International Color Consortium has agreed on a document which describes a flexible and extensible file format for storing device profiles. This is something of an achievement as it follows failed efforts by other organizations to standardize different aspects of color management. However, it must be recognized that the file format is only one part of the software infrastructure necessary to manage color. Since the technical language is still evolving, it is worthwhile to describe the parts of this structure.

In the jargon of color management, an input (or source) *profile* provides the information required to interpret the digital code values of a color image or graphic. An output (or destination) profile allows the completion of the communication process by providing the information needed to reproduce the colors on an output device. The software that uses the information contained in profiles to accomplish this translation is called a *color management module* (CMM). On most computer platforms, the operating system will provide software *application* developers access to these capabilities through a color management *framework* (and its associated application programmer's interface or API). Most frameworks will allow end-users or applications to replace the *default CMM* included with the OS by a third-party CMM. Nevertheless, there will inevitably be applications which use ICC profiles but which do not work through the OS framework.

There are thus four key components of the ICC system: the profiles, the CMMs, the frameworks, and the applications. The ICC format is intended to serve as a de facto standard which allows these pieces to function as an "open" system. If profiles interoperate, users can choose the components from different vendors yet still expect a basic level of color management. Facilitating interoperability was the most important goal of the consortium and it had considerable influence on how color is managed in the ICC system.

The Colorimetric Context

Interoperability makes it easier to communicate color images and graphics between different users. This requirement is added to the other task of color management: to communicate color between devices and viewing environments.

For our purposes, managing color may be defined as a process of translating its digital representation so as to facilitate its reproduction. The goal of color reproduction is essentially the communication of color appearance. So a color management system must find a way to quantify how an image appears in order to communicate it and ultimately allow its reproduction. This is achieved through the definition of what has been termed a *reference color space*. The reference color space is usually labeled "device-independent" since it is based on CIE colorimetry rather than the properties of a specific imaging device.²

As stated above, the device profiles allow the interpretation of device-specific digital codes by establishing their relation to the reference color space. However, since the color reproduction process is a translation from one device space to another, the reference space is *virtual*. Images need never reside in this space nor must it have a specific encoding with the associated range limits and quantization. Essentially, it is only an abstraction used to connect profiles. It is therefore more appropriately called a *profile connection space* or PCS.

Since colorimetry only describes the physical stimulus and not the perception it evokes, there is no unique color appearance associated with a set of three tristimulus values (or CIELAB coordinates, etc.). For example, identical tristimulus values appear differently as seen in the context of a reflection print, computer monitor, or transparency viewer. Stated from an alternate point of view, colorimetry must be adjusted to preserve the same appearance across different media and different viewing environments.³ There must therefore be an associated *colorimetric context* from which color appearance can be inferred.

The colorimetric context must describe both the nature of the imaging media and the viewing environment. However, some of this information may be implicitly agreed upon and fixed when the PCS is designed, and some of it may be communicated explicitly along with the associated transforms. The more information that is fixed, the more *narrowly* the PCS is defined. If most of the context is not fixed a priori, then the definition of the PCS is considered to be *broad*. As will be explained, there are advantages to both approaches.

We should draw a distinction between appearance as communicated through a PCS from appearance as defined by a model such as the Hunt color appearance model.⁴ The Hunt model defines variables which correlate with perceptual quantities such as brightness, lightness, colorfulness, and hue. The variables used to encode the PCS could also be defined in these terms; however, this is not a requirement. A PCS communicates color by specifying the colorimetry required to *reproduce* the desired color appearance in a specific situation defined by the colorimetric context. In fact, a given color management system may simultaneously support several different PCS encodings without changing the reproduction of colors.

A PCS has features in common with the RLAB space proposed by Fairchild and Berns.⁵ The RLAB space is notable for its emphasis on the need for accommodating differences across viewing environments. However, it does not deal with the media related issues we discuss below. That is, it defines a viewing context but not a media context.

Defining the Viewing Environment

The part of the colorimetric context which is dependent upon the viewing environment is due to the complex, adaptive nature of human vision. To the extent that these factors can be fixed during the design of the PCS, a *reference viewing environment* is defined. The RLAB space

discussed above uses the concept of a reference viewing environment. Those factors which are not fixed as part of a reference must be communicated as part of the profile. We have divided our discussion of the factors that define a viewing environment into three areas: adaptive effects, inductive effects, and veiling glare.

Adaptive Effects

The absolute luminance of a diffuse white ranges from very roughly 20,000 cd/m² for a sunlit scene, to 700 cd/m² for a print in an ANSI standard viewing booth, to 100 cd/m² for a computer monitor, to even lower values for a slide or movie projected in a dark room. The chromaticity of neutrals varies (as expressed by the correlated color temperature) from about 2000K for tungsten illuminated prints to over 10,000K for some computer displays. The human visual system uses physiological and cognitive methods to reduce the effect of changes in both the level and chromaticity of illumination on the appearance of a scene; the process is referred to as *adaptation*.⁶ Since the process is not complete, sophisticated models of adaptation will correct colorimetry for changes in contrast as the absolute luminance level changes and modify the relative hues of colors as the color temperature of white changes.^{7,8} More typically, the actual colorimetry is simply normalized to a reference white (though this is often unsatisfactory).

Because of the sophisticated cognitive aspects of adaptation, modeling it requires more information about the viewing environment than the absolute tristimulus values of a reference white. For example, adaptation is more complete in a dark room than a light one. Also, an observer is more likely to *discount* (or completely adapt to) the color of a neutral on a CRT than on a print even if both are viewed in the same office. Even the type of graphical user interface (GUI) used in a computer display can affect adaptive mechanisms.

Inductive Effects

The *surround* typically refers to the area of the visual field that extends from the edge of an image or target in all directions. Often it is not uniform, such as in many computer displays or reflection print viewing environments. Sometimes the surround consists of a whitish border, itself surrounded by a darker color. In any case, the effect of the surround is to induce a complementary sensation in the area surrounded.⁹ Thus, an image with a dark surround appears lighter than the same image with a light surround. The amount of the effect is dependent not only on the level and color of the surround but also upon the visual angles subtended by both surround and image. Obviously, the tone and color reproduction of an image needs to be modified for the type of surround expected if appearance is to be communicated. The exact nature of the correction is a matter of some debate but often consists of a power-law type transform applied to the tristimulus values.

There are other effects which may also be loosely referred to as inductive effects, such as assimilation and simultaneous contrast. However, these are dependent on the content of an individual image. While, by definition,

the colorimetry of the image pixels affects the appearance, the pixels should not be corrected for such intra-image effects. On a pixel level, colorimetry essentially loses its value—indeed, the environment created by the surrounding pixels violates the assumptions upon which color matching experiments rely.

Veiling Glare

For reasons too numerous to describe here, the images we see are contaminated by stray light. Also known as *flare*, this causes dark areas and saturated colors in an image to lighten or “wash out” while leaving whiter areas relatively unchanged. It is typically specified as a percentage of the intensity of a reference white. Almost all practical viewing environments have at least a half percent, but some environments have much more (ex. a CRT in a typical office). The nature of the flare is sometimes dependent on the image content, although in practice it is usually considered to occur in a non-image-wise fashion.

While adaptation and induction concern the effects of the environment on the visual system, flare concerns the effect of the environment on the physical stimulus itself. That is, the colorimetry that would be measured by a (typically more convenient) flareless measuring process differs from the physical stimulus viewed in situ. Other environmental factors being equal, the colorimetry *as viewed* in the receiver’s environment should be the same as that viewed in the sender’s environment regardless of the difference in flare. In this sense it is different from adaptive and inductive effects in which the colorimetry viewed must change to keep the appearance the same.

Communication requires agreement on the sort of colorimetry being sent, in this case, whether or not it includes viewing flare. If it does not, then a correction should be applied for the differences in the viewing environment. If it does, then no correction is needed for the difference in viewing environments, but it remains necessary to correct colorimetry measured without flare so that it reflects the actual viewing conditions. A simplistic correction may be achieved by applying a small offset to the tristimulus values.

Defining the Media

The part of the colorimetric context that is dependent upon the imaging medium is due to the limited extent of the medium’s color gamut. (We use the term “medium” here to refer to the combination of an imaging device acting upon materials such as paper, dyes, inks, or phosphors.) No single medium contains the entire gamut of colors found in all scenes which someone may want to reproduce—even if the “scenes” are limited to reproductions on typical input media. Hence a color management system is faced with the problem referred to as gamut mapping, or perhaps more specifically, *cross-media mapping*. By using cross-media mapping, we wish to emphasize that it is not solely a function of the output medium—the input medium is involved as well.

Without some knowledge about both the source and destination media, the cross-media mapping problem is intractable. For example, the source may have the

100,000:1 dynamic range of a scene, the 4000:1 range of a transparency, or the 100:1 range of a magazine and each would require quite a different mapping for any given output. To do the best possible job at cross-media mapping requires knowledge of the specific media used for the source and destination reproductions. As an example, one might construct a custom mapping of the Ektachrome transparency gamut onto the Matchprint proofing system gamut (and another for Kodachrome onto Matchprint, etc.). However, color management has not evolved to the level at which this sort of device-to-device-specific mapping is practical.

In practice, color management systems will use very specific knowledge of the output gamut but only general knowledge of the input gamut. For example, an output profile might map a transparency type of product onto the Matchprint gamut. If the PCS may also contain reflection print type inputs or “scene” type inputs, then additional output profiles may be needed to map these inputs onto the Matchprint gamut as well. Otherwise, the CMM might have to apply a mapping to accommodate an input not intended for a specific output profile. However, if the PCS only represented colors as rendered on one type of media, one reduces the interdependence of input and output profiles. While the problem of doing gamut mappings in output transforms is widely discussed, the corresponding issue concerning the making of input profiles is largely ignored.

As a further example, consider a PCS to SWOP-press CMYK transform in an output profile that “compresses” the colors in the PCS to fit into the press gamut. There should be a corresponding input transform that allows the SWOP-press-targeted CMYK data to be brought into the PCS as well. The profile builder may decide to simply transform the measured colorimetry into the PCS. However, this would not actually invert data from the output transform—there would be a residual gamut compression. If this data were to be sent to another device through another compressing output transform the data would be compressed twice. The profile builder may alternatively undo the gamut map (of the corresponding output transform) in the input transform. This would then allow a “round trip” to be more loss-less. Both of these approaches to making input profiles can be useful in facilitating different types of reproduction. However, it is obviously important that the mapping in input and output profiles be coordinated.

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. In some sense, these fixed properties define what might be called a *reference medium*. Transforming colors into the PCS then involves replacing the characteristics of the source medium with those of the reference medium. On output, the characteristics of the reference medium are replaced by those of the specific destination medium. The advantage of using a PCS with associated media properties 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.

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. Not associating media properties with the PCS makes it easier to render the 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 good example of this type of application would be the use of a dye-sublimation printer to produce an image which is a colorimetric copy of how that image would look if printed on a SWOP-press. In this case, the reproduction is not intended to be the most “pleasing” reproduction possible on the dye-sub media but rather a simulation of another medium. While this application can also be accomplished using a PCS with an associated reference medium, it requires using a “pre-view” or “simulation” transform.

A final point to make regarding the media related properties of the PCS is the following: defining PCS media properties broadly is different from not defining them at all. A broadly defined PCS either has no associated reference medium, or may only define certain aspects of the reference, or define it in a loose or general sense. There is a continuum between narrow and broad definitions. However, those aspects which are not fixed should be discernible from the data fields in the profile.

Profile and CMM Interaction

Color appearance is communicated through interpretation of the colorimetric context associated with a PCS. Therefore, 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.

If the PCS context is variable, it will be communicated along with a profile. Recall that the CMM performs source to destination color translations by combining a series of profiles. If the interpretation of the PCS is not the same in all of the profiles involved in such a translation, it is up to the CMM to perform the necessary accommodation. This accommodation will require performing cross-media and cross-environment mappings—making the CMM a rather sophisticated and complex bit of software. The profiles would contain the data needed by the CMM and defining the variable aspects of the colorimetric context. In this scenario, the profiles might contain transforms that treat PCS values as simple colorimetry.

Perhaps the main disadvantage of doing the mappings in the CMM is that they would have to be done very quickly. A CMM can not pre-compute these mappings since it will not have prior access to the input profiles. Although a CMM will typically have access to all of the output profiles in use on a given system, input profiles will be arriving along with color documents. The profile specification describes the manner in which profiles may be enclosed in TIFF files (embedding procedures

for PICT and PostScript files are also being addressed). 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.

If the PCS context is fixed, then the CMM need not address the problem of whether appearance is being communicated properly from one profile to another. While this simplifies the CMM, it requires that the cross-media and cross-environment mappings be incorporated into the profile building process. This means that the transforms in a profile relate device codes to colorimetry *as adjusted for the fixed colorimetric context*.

The PCS and Interoperability

The choice of where the cross-media and cross-environment adjustments are made has system implications. For example, the means of updating the installed base of users as technology improves depends dramatically on this decision. More importantly, the choice affects the relative difficulty of building profiles versus building CMMs and the type of interoperability facilitated (cross-profile vs. cross-CMM).

The differences in media and viewing environments between source and destination will often preclude a complete reproduction of the color appearance. This forces the cross-media and cross-environment mappings to be at least partly based on empirical methods which give an aesthetically successful, or “pleasing,” compromise. These methods are often proprietary. Furthermore, there are inter-relationships: the effective media gamut is altered by changes in the viewing environment, and the viewing environment corrections are limited by the edges of the gamut. As a result, it is not currently possible to standardize these mappings between vendors. It is quite difficult to agree upon the lesser goal of what factors define the colorimetric context. Even distinguishing between environment and media related aspects of an overall color translation is difficult.

If cross-media and cross-environment mappings will be somewhat different as implemented by different vendors, the above design choices have the following effects on interoperability. A simple CMM design that does no more than assemble profiles increases interoperability of a given profile between different operating systems, computer platforms, and applications. Simplifying profile building (by forcing the mapping into CMMs) facilitates interoperability of profiles from different vendors on a given vendor’s CMM. Also, regardless of where they are located, interoperability of profiles and CMMs will be improved if both media and environment mappings are done by the same vendor (hence in the same component).

The ICC Profile Specification

This paper has described the broad color communication issues which we feel influenced the design of the ICC Profile Format. Hopefully, this will serve as useful background information for those examining the Profile Format in detail. In this section, we comment on several

design choices made by the Consortium and described in the specification document.

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. The CMM need only convert between a series of several encodings of the PCS. Allowing a simple CMM design should greatly facilitate the interoperability of ICC CMMs from different vendors. In special situations where more sophisticated CMMs are needed, the advanced CMM should provide a “compatible” or “default” mode of operation where any alteration of PCS colors is turned off.

The implied simplicity of the CMM requires profile builders to include any cross-environment and cross-media mappings deemed necessary into the transforms themselves. The specification goes into some detail describing what these mappings should accomplish. Hopefully, this will increase the interoperability of ICC profiles from different suppliers.

One potentially confusing aspect of the specification concerns the relationship between the ICC PCS and a set of three colorimetric encodings mentioned (a 16 bit/channel XYZ, an 8 bit/channel CIELAB, and a 16 bit/channel CIELAB). These encodings simply define the representation of the 8 and 16 bit color transforms. The PCS itself does not have specific quantization, encodings, or range limits. For example, although the two CIELAB encodings have an upper limit on L^* of 100 it is conceivable that colors in the PCS may exceed this limit.

Further confusion may be caused by discussion in the specification of *absolute colorimetry* and *relative colorimetry*. These do not have direct bearing on the definition of the PCS: they are used to provide different types of gamut mappings (as described in Appendix D). The transform for the absolute rendering style is not provided explicitly in the profile format. Rather, provision is made for the CMM to generate it at run-time through the use of a tag describing the media white point.

The so-called absolute colorimetry is really just standard CIE colorimetry.¹⁰ That is, the tristimulus values of colors are normalized by the tristimulus values of a perfectly diffusing white reflector. Rendering styles which reproduce CIE colorimetry are useful when trying to achieve an exact colorimetric match between two different media, as when printing “spot colors.” However, the absolute rendering style can result in unpleasant color reproduction.

One problem associated with matching absolute colorimetry is poor rendering of whites. Since the density of different media “whites” vary, image detail near white may be clipped off or rendered with an unpleasant cast depending on the media involved. A convenient, while perhaps not optimal, solution is the use of the so-called relative colorimetric rendering style. This is based on what might be called *media-relative colorimetry* since instead of normalizing the tristimulus values by those of an ideal white reflector, a different normalization is chosen based on the media and viewing environment in use. This enables profiles to be plugged together in a conve-

nient way while ensuring that white objects on input properly map to white objects on output. The use of media-relative colorimetry is an example of how media properties are incorporated into the colorimetric context of the PCS, as discussed earlier.

Conclusion

Realistically, it will take some time before the companies involved in implementing the ICC infra-structure completely standardize on the solutions to the given communication problems. Therefore, the highest quality results will initially be obtained by using profiles and CMMs designed to work together. If profile vendors take radically different approaches, users will have to use profiles from a single vendor for best results. If CMMs attempt to apply their own mappings when combining profiles, users communicating between operating systems or applications will require more than one profile for a single device (depending on the CMMs being used).

As the reader may have concluded, there is quite a lot involved in correctly building an ICC profile. Fortunately for application developers, profiles will either be generated through CMM calls which hide the complexity, or be available as pre-constructed entities from specialized profile generation applications or profile vendors. For color scientists, the implementation of the ICC infra-structure will require consensus building on what constitutes the proper cross-media and cross-environment transformations. Since there are questions of aesthetics involved, these transforms will never be identical as realized by different vendors, but they must at least be interoperable. If this goal can be achieved, users will gain unprecedented flexibility and power to communicate with color images.

References

1. The ICC Profile Format specification will be available from FOGRA, a European graphic arts organization and via ftp. Details are not set as of this writing. We can provide current information on availability; you may send email to walker@keps.com.
2. For example, D. Walker, Color management—three advanced topics, in *Proc. of IS&T's Third Technical Symposium on Prepress, Proofing & Printing*, Chicago, 1993.
3. For example, R. W. G. Hunt, *The Reproduction of Color in Photography, Printing, and Television*, 4th ed., Fountain Press, England, 188-190, 1987.
4. R. W. G. Hunt, Revised colour-appearance model for related and unrelated colours, *Color Res. Appl.*, **16** (1991).
5. M. D. Fairchild and R. S. Berns, Image color-appearance specification through extension of CIELAB, *Color Res. Appl.*, **18** (1993).
6. For example, L. M. Hurvich, *Color Vision*, Sinauer Assoc., 1981.
7. Hunt, op. cit.
8. Y. Nayatani, K. Takahama, and H. Sobagaki, Prediction of color appearance under various adapting conditions, *Color Res. Appl.*, **11** (1986).
9. For example, Hunt, op. cit.
10. CIE Publication No. 15.2, *Colorimetry*, 2nd ed. (1986).