The History of Device Independent Color – 10 Years Later

Robert R. Buckley
Solutions & Services Technology Center
Xerox Corporation, Webster, New York

Abstract

The First Color Imaging Conference in 1993 was organized around the theme of “Transforms and Transportability” and included an invited paper entitled “A Short History of Device Independent Color.” Several developments since then have affected the idea of device independent color and its application to color imaging. These developments include the ICC, which was founded in the same year that this conference series started; the formation of CIE Division 8; the introduction of sRGB and other RGB color spaces; the standardization of color appearance models; the growing use of spectral reproduction techniques; the development of the image state concept; and the phenomenal growth of the web. This paper will update the history of device independent color, describing how device independent color and color interchange and reproduction have been influenced by these developments and what we’ve heard at past Color Imaging Conferences.

Introduction to Device Independence

Although the notion of device independent color originated in the 1980's, the idea of device independence itself goes back a decade or two earlier in the field of computer graphics. The goal in computer graphics were applications that were not tied to a specific output device and its characteristics, but could generate equivalent-looking results on a variety of display devices, independent of whether they were vector or raster, and if they were raster, what their spatial resolution was.

This goal was achieved by using a virtual or reference coordinate system to describe the desired appearance of the output. This approach decoupled the way in which an application described an image from the way in which a device rendered it. For example, an application could use a reference system for spatial coordinates with its origin at the lower left-hand corner of a page and meters as units. The coordinate system for a raster output device could put the origin at the upper left-hand corner of the page and use pixels as units. The transformation between reference R and device D coordinates is then a linear transformation T_2 (Figure 1), whose coefficients depend on the resolution and orientation of the reference and output device coordinate systems.

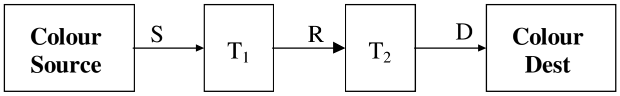

Figure 1. Transforms between color source and destination devices

This approach to device independence can be extended to include input devices by means of a transformation T_1 from input device coordinates S to reference coordinates R.

In a device-independent system based on reference coordinates, all input devices, when combined with the transformation from their device-dependent input coordinates to reference coordinates, look like a standard input device, a virtual device that uses the reference coordinate system. And all output devices, when combined with the transformation from reference coordinates to their device-dependent output coordinates look like a standard output device, a virtual device that uses the standard reference coordinate system. As far as an output device is concerned then, all input devices look the same, or at least have the same interface, and vice versa.

This approach to device independence, or more precisely dependence on a standard reference device, has advantages for both the user and the implementer. The user (client or content designer) can create and specify content without having to know the details of how the output device renders content. The implementer can build to well-defined interfaces, increasing the opportunity to share and reuse pieces of the implementation that depend on the reference coordinates. In fact in graphic systems, device independence has been used to describe “a system that allows programs to be used unchanged with different input and output devices”.[1] In other words, device independence originally was as much if not more concerned with the portability of implementations as it was with the portability of data.

This discussion of device independence has so far only considered the orientation and units of the spatial coordinate system. In a practical, ultimately device-oriented system there are two other aspects of coordinate systems to be considered: the range and the resolution (or addressability) of the coordinate values. These two properties are determined by the encoding of reference coordinates and by the physical dimensions of the device coordinate space.

For example, using floating point numbers to encode coordinates values does not practically limit the range or resolution of coordinate values. Integer coordinate values on the other hand may limit both. However, using A4 paper for either input or output would limit the range of coordinate values, even if the encoding itself didn't. In effect, the range of coordinate values that occur in a device-independent encoding are device dependent. In the case, A4 paper defines the spatial gamut of the device and limits the range of device-independent coordinate values.

Device Independent Color

The obvious choice for device-independent or reference color coordinates are the CIE XYZ tristimulus values for the 2° Standard Observer, or a CIE-specified transform of them, such as CIE L*a*b* values. In these cases, the virtual device is a colorimeter or equivalently the 1931 Standard Observer, which provides a standard and objective measure of a color stimulus that applies to input and output devices. It is well known however that CIE values are insufficient to describe the appearance of color stimuli. Except under constrained conditions, CIE values by themselves cannot lead to the equivalent-looking results that device independence seeks to achieve across output devices. This is the added challenge of color device independence, beyond the limits on device independence noted earlier in the spatial case. A meter stick is always measures out to be a meter long in normal situations, but a D50 color stimulus does not always look white.

Besides XYZ tristimulus values, the appearance of a color stimulus depends on the conditions under which it is viewed, which are effectively part of the device definition. A viewing-conditions-independent color representation, based on a color appearance model, would come closer to device independent color [Ref. 2, p.346]. Research in color appearance modeling reached the point where the CIE published the CIECAM97s interim model in 1998.[3,4] Work has continued on color appearance modeling.[5] and a revised version of the CIECAM97s model called CIECAM02 has been developed.[6]

From time to time, color appearance models have been brought up in connection with color management. However, existing models have been too complex, their parameters not always known, and their applicability to images limited for them to see significant use in commercial color imaging systems. So device independent color is still based on CIE tristimulus values, with viewing conditions carried as metadata or made part of the reference or virtual device's definition.

The computation of CIE L*a*b* values includes a normalization to a reference white, which allows for a crude white-point-independent color representation. A change in the reference white changes the color appearance—a change usually compensated for by a chromatic adaptation transform—but more than that it changes the tristimulus values of a hardcopy color stimulus. In the latter case, using the spectral reflectance curve—or a low-dimensional approximation to it—provides a more illuminant- and therefore more device-independent color representation than tristimulus values.[7] Further, a multi-spectral stimulus representation enables more robust image reproduction than a tristimulus-value-based approach[8] and is being applied to the image archiving and reproduction of fine art.

One aspect of CIE XYZ tristimulus values and CIE L*a*b* values that make them suitable as device-independent color coordinates is that they cover the range of all physically-realizable color stimuli. To provide sufficient resolution (fine enough quantization) in a byte-oriented digital system, the preferred encoding of XYZ values would use 16 bits. With CIE L*a*b* values, which are a non-linear transform of XYZ values and a better match to color appearance than XYZ, 8 bits usually give sufficient resolution. However, 8 bits limit the a* and b* values of color stimuli that can be represented to the range [-128, 127]. Although there are physically-realizable colors outside this range, typical hardcopy colors are within it.

An important set of transforms of CIE XYZ values are the ones that convert to RGB values, including monitor RGB values. An XYZ-to-RGB transform itself does not restrict the ability of the resulting RGB coordinate values to represent all physically-realizable colors. However, the encoding typically limits the set of addressable colors to just those within the monitor gamut. For example, sRGB[9,10] uses an 8-bit non-linear encoding with non-negative integers that span the gamut of a monitor with CCIR 709 primaries, so that the encoding is device dependent. Other encodings of sRGB, such as scRGB[11] and e-sRGB,[12] allow for negative coordinate values and extend the range of colors that can be represented.

The other approach to increasing the coverage of an RGB color encoding while retraining an encoding with non-negative integers is to move the primaries out to obtain wide-gamut RGB color spaces.[13] Examples of wide-gamut RGB color encodings are ROMM-RGB[14] and RIMM-RGB,[15] which uses the same primaries as ROMM-RGB but also allows an extended range. All of these extended encodings support more than 8 bits.

Some "device" of course determines the values that actually occur and need to be represented, whether the encoding is device independent or not. It could be the output device for which the values are targeted, or the input device used to capture a scene or original. The notion of image state has been introduced to describe the nature of color values an encoding was designed to represent.[16] Figure 2, based on Figure 1 in Ref. 17, shows input- and output-referred encodings, related by a transform T. An output-referred encoding can be targeted at a printer or monitor. Input encodings can be either original-referred, in the case of the scan of a print or slide, or scene-referred, in the case of values from a digital camera, a color film negative scanner or a computer-generated scene. Each of these represents a device's rendering of color data, so that image state is device dependent.

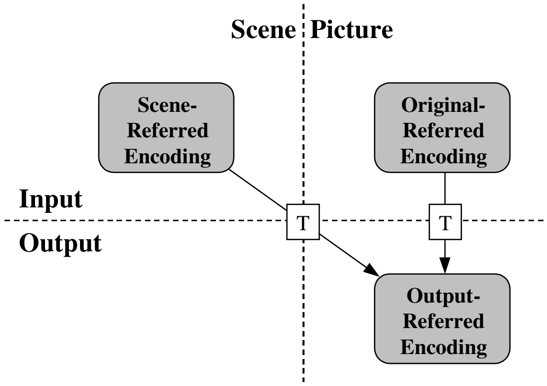

Figure 2. Image states and encodings

Interchange Models and Transforms

In a connected system of input and output color devices, the key to successfully achieving color data interchange and color reproduction is the use of transforms between device dependent and device independent color representations. Figures 1 and 2 showed the conceptual partitioning of devices and transforms to create virtual devices. By comparison, Figure 3 shows the physical configurations of color sources and destinations and their associated transforms in a color system.[18,19]

The choice of interchange coordinates defines the external interfaces in these configurations and determines system characteristics and performance. The attributes that differentiate these configurations are:

- Bandwidth use on the interchange channel
- Ability to use different sources or destinations
- Distribution of processing capability and cost among system elements
- Compatibility with archived color data or existing device capabilities

The configurations in Figure 3 differ in terms of the coordinates they use for color data interchange:

(a) destination-targeted (output-referred) device coordinates D;
(b) device-independent reference coordinates; or
(c) source-defined (input-referred) device coordinates S.

Destination-Targeted Interchange

Configuration (a) emphasizes simple color processing at the destination, with the source exercising most of the color control and doing most of the processing. For the source to export destination-targeted coordinates D, it must have knowledge of the destination, either explicitly using side information from a back channel or implicitly by making assumptions about the destination. For example, this configuration may assume an output device that performs according to published specifications. Although the transform at the source is shown in Figure 3(a) as two transforms, there is no particular requirement on how the transforms are implemented and configured at the source. The source could, for example, use a single, merged transform based on ICC DeviceLink profile data.

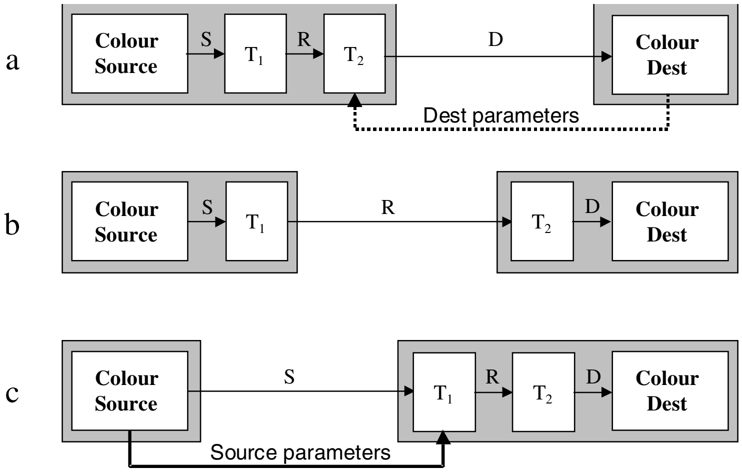

Figure 3. System configurations for color data interchange

Examples of configuration (a) are systems that use SWOP CMYK or sRGB. From the standpoint of color, another example is the NTSC color television system. The NTSC system was designed to operate with a receiver that had a specific phosphor set and transfer characteristic, and produced RGB values for that specific receiver. What were actually broadcast were luminance-chrominance values, but the overall effect of converting to RGB to luminance chrominance at the transmitter for broadcast and then back to RGB at a color receiver was the identity transform.* In many applications, the transformation from RGB to luminance-chrominance and back again to RGB is invisible to the user and is performed to improve the compressibility of an RGB color image, using JPEG or JPEG 2000 for digital image compression for example.Accurate color on the web is a difficult problem.[20] The desire to use the web in applications, such as catalog sales, where accurate color is important has led to an interesting example of configuration (a) for viewing images in a web browser. A simple, interactive procedure is used to characterize a client's monitor. The results are sent to a server where they are used to derive a transform, which the server applies to the catalog images, correcting them for the client's monitor. When a user downloads a catalog web page, the references on the web page then would point to these corrected images on the server. E-Color's True Internet Color® architecture is an example of this configuration. A variation is Gretag-Macbeth's WebSync, which an applet with the uncorrected image; the applet color corrects the image in the client browser. Whether the color transformation is done at the server or at the client, it is based on the characterization the client's monitor. Moving it to the client is an example of configuration (b).

Device Independent Interchange

Configuration (b) in Figure 3 emphasizes the interchange channel and decoupling the source and destination. The source explicitly transforms source coordinates to standardized, device independent coordinates, which the destination transforms to the device-dependent coordinates appropriate for it. This configuration is symmetric and relatively balanced in the distribution of processing between source and destination.

A color facsimile system is an example of configuration (b), where the reference coordinates are CIELAB, assuming a D50 illuminant and with the range of the a* and b* coordinates scaled so that 8 bits tightly cover the range of values found in hardcopy.[21,22] Another example is the use of ROMM-RGB encoding as the reference coordinates.[14] It is also possible for Configuration (a) to migrate Configuration (b) when the destination device is different from the "standard" one for which the source prepared device dependent coordinates D. In this case, a transform inserted at the destination re-targets the standard, device-dependent coordinates to the ones suitable for the actual destination device.

Source-Defined Interchange

In configuration (c), the emphasis is on flexibility in source-destination color pairings, where the source needs little or no knowledge of the destination's color capabilities. Color transforms are deferred until it is time to render the color at the destination, i.e. late binding. In this configuration, the source exports its coordinates along with the parameters that are used to construct the transform from source to reference coordinates.

The use of the calGray and calRGB color spaces in Adobe's Portable Document Format (PDF) are examples of parameterized transforms from source gray or RGB values to CIE XYZ reference values. These color spaces first appeared in 1994 in the PDF 1.1 specification. They are instances of the CIE-based color space model, which is supported by the PostScript language.

ICC color management is also an example of Configuration (c).[23] The ICC or International Color Consortium was founded in 1993, and had its origins in Apple ColorSync, which was described at the First Color Imaging Conference.[24]

In an ICC-based system, the parameters of the source-to-reference transform are contained in an input or display device profile, which follows the ICC Profile Format standard.[25] The transform can take one of a few forms and is constructed from a set of "standard" components or processing elements. The reference coordinates are defined by the Profile Connection Space (PCS), which can use either CIE XYZ or L*a*b* values. Similarly, the transform from the PCS to destination device coordinates can also be described in terms of an ICC output profile. At the destination, these two profiles or transforms are combined in a Color Management Module (CMM) to obtain the overall transform from source to destination color coordinates.

When the time comes to render the color and the source and destination color gamuts differ, a gamut mapping is needed. For best results, the gamut mapping would be different for different kinds of content and their associated intent. In some cases, it is best to render the colors as closely as possible to their colorimetric specification. In other cases, such as rendering pictures, all colors should be adjusted so that the relationships between them are roughly the same after gamut mapping as they were before, in accordance with Evan's Consistency Principle. In the case of charts or graphics, the colors should be rendered as saturated as possible. For example, when a yellow is specified, the intent is to have a clean, solid yellow color rather than one that is mixed with others colorants even though it might be closer to the specified colorimetric value.

The ICC captures these three gamut mapping strategies with the notion of rendering intent: colorimetric (absolute and relative), perceptual and saturation. Gamut mapping algorithms are beyond the scope of the ICC standard, but are being investigated by CIE TC8-03, which is one of the

---

[*] This ignores the effect of filtering the luminance and chrominance signals at the receiver.

technical committees of CIE Division 8, created in 1998. The current state of the Division's work is reported in Ref. 26.

Besides providing reference coordinates, the PCS also now defines a reference device, at least in the case of the perceptual transforms. This definition of the reference "device" includes a reference medium whose white and black points have reflectances of 89% and 0.30911%, and reference viewing conditions that follow ISO 3664. A procedure for converting source device coordinates into PCS reference device coordinates in the case of the perceptual intent is outlined in Annex D.2.7.8 of Ref. 24. In this case, device independence is, as was noted earlier, dependence on a reference device. All that's left that can be device independent are the coordinates, and then only when they don't limit the range of color stimuli that can be represented.[27]

At each of the source, reference and destination devices, the color stimuli exists in a rendered state. Simply converting the coordinates between states without regard to viewing conditions and other metadata can be considered "transcoding." If the viewing conditions are the same in the two states, then transcoding is sufficient. An example of transcoding is converting from PCS coordinates to destination CMYK values when the reference device describes the destination device. RGB-to-YCbCr conversions are also an example of transcoding. A more general state transform is "transrendering," where the "devices" differ in ways other than the coordinates they use. An example of transrendering is converting from softcopy source coordinates, which typically use a D65 white point, to PCS values. Transcoding and transrendering are two types of the transform that have been and remain central to the notion of device independent color.

Conclusion

This update to the history of device independent color has had two main objectives: to show how the idea and practice of color interchange has evolved and advanced since this conference series began; and to show that, when you come right down to it, device independent color is not literally achievable. The latter may seem like a trivial observation, but the implications are significant for those who actually need to make color interchange and color reproduction work satisfactorily. More than that, users will want to use the interchange methods described here to address the specific capabilities of a preferred output device.

What will the next update to the history of device independent color look like? It will undoubtedly draw on progress in the International Color Consortium and in CIE Division 8. This new CIE Division, with its focus on Image Technology and including technical committees on color appearance, gamut mapping, color communication and multispectral imaging, is investigating the technical issues associated with color imaging. There are also the standard committees that make technical results available to practitioners in agreed-upon ways. Let us hope that this conference series will continue to attract and bring together all those working and contributing to color image technology, and to the interchange of ideas as well as color images.

References

1. W.M. Newman and R.F Sproull, *Principles of Interactive Computer Graphics*, 2nd Edition, McGraw-Hill, New York, 1979, p. 109
2. M. Fairchild, *Color Appearance Models*, Addison-Wesley, Reading, 1998
3. CIE Publication 131, The CIE 1997 Interim Colour Appearance Model (Simple Version) CIECAM97s, Commission Internationale de l'Eclairage, Vienna, Austria (1998)
4. A.R. Robertson, J. Schanda and T. Newman, CIE Work in Color Imaging, *Proc. Sixth IS&T/SID Color Imaging Conference,* Scottsdale, Arizona, November 1998, pp. 14-16
5. M.D. Fairchild, Status of CIE Colour Appearance Models, *Proc. Ninth Congress of the International Colour Association*, Rochester, NY, June 24-29, 2001, pp. 550-553
6. N. Moroney, M.D. Fairchild, R.W.G. Hunt, C. Li, R. Luo and T. Newman, *The CIECAM02 Color Appearance Model*, *Proc. Tenth IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 2002, pp. 23-27
7. T. Keusen and W. Praefcke, Multispectral Color System with an Encoding Format Compatible to the Conventional Tristimulus Model, *Proc. Third IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 1995, p. 112
8. P.D. Burns and R.S. Berns, Analysis Multispectral Image Capture, *Proc. Fourth IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 1996, pp. 19-22
9. M. Anderson, R. Motta, S. Chandrasekar and M. Stokes, Proposal for a Standard Default Color Space for the Internet—sRGB, *Proc. Fourth IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 1996, p. 238
10. IEC 61966-2-1: Multimedia systems and equipment – Colour measurement and management – Part 2-1: Colour management – Default RGB colour space - sRGB
11. Draft IEC 61966-2-2: Multimedia systems and equipment – Colour measurement and management – Part 2-2: Colour management – Extended RGB colour space - scRGB
12. PIMA 7667:2001, Photography – Electronic still picture imaging system – Extended sRGB color encoding – e-sRGB
13. S. Süsstrunk, R. Buckley and S. Swen, Standard RGB Color Spaces, *Proc. Seventh IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 1999, pp. 127-134
14. PIMA 7666:2001, Photography – Electronic still picture imaging system – Reference Output Medium Metric RGB Color Encoding: ROMM-RGB
15. I3A 7466, Photography – Electronic still picture imaging system – Reference Input Medium Metric RGB Color Encoding: RIMM-RGB
16. K.E. Spaulding, G.J. Woolfe and E.J. Giorgianni, Image States and Standard Color Encodings, *Proc. Eighth IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 2000, pp. 288-294

17. ISO TC42, ISO/CD 22028-1, Photographic and graphic technology – Extended colour encodings for digital image storage, manipulation and interchange – Part 1: Architecture and requirements, April 4, 2002
18. R.R. Buckley, A Short History of Device-Independent Color, *Proc. First IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 1993, pp. 88-92
19. R.R. Buckley, Device Independent Colour and Virtual Colour Devices, *Proc. International Colour Management Forum*, University of Derby, March 1999, pp. 71-78
20. R.W.G. Hunt, How to Shop on the Web without Seeing Red, *Proc. Eighth IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 2000
21. A. Mutz and D.T. Lee, An International Standard for Color Facsimile, *Proc. Second IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 1994, pp.52-54
22. D.T. Lee, Introduction to Color Facsimile: Hardware, Software and Standards, *SPIE* Vol. **2658**, 1996, p. 8; reprinted in [28]
23. J. King, Why Color Management, *Proc. Ninth Congress of the International Colour Association*, Rochester, NY, June 24-29, 2001, pp. 439-445
24. G. Murch, Color Management on the Desktop, *Proc. First IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November 1993, pp.95-99
25. International Color Consortium, File Format for Color Profiles (Version 4.0.0), ICC.1:2001-12, www.color.org
26. T. Newman, CIE Division 8 and the Philosopher's Stone, *Proc. Tenth IS&T/SID Color Imaging Conference*, Scottsdale, Arizona, November, 2002, pp. 47-50
27. E. J. Giorgianni and T. E. Madden, *Digital Color Management*, Addison Wesley, Reading 1998
28. R.R. Buckley, ed., *Recent Progress in Color Management and Communications*, Society for Imaging Science and Technology, Springfield, VA, 1998

Biography

Robert Buckley is a Research Fellow and the Manager of Color & Imaging Solutions in the Xerox Solutions and Services Technology Center. He received a BScEE from the University of New Brunswick, an MA in Psychology and Physiology from the University of Oxford and a PhD in Electrical Engineering from MIT. He is the Xerox Principal on the JPEG 2000 Committee and Editor of Part 6 of the JPEG 2000 standard, and he is the Chair of CIE TC8-05 on the "Communication of Colour Information." He is also a member of the Board of the Inter-Society Color Council.