Specification of Colorspace in JPEG 2000 Balances Flexibility with Interoperability

J. Scott Houchin
Eastman Kodak Company
Rochester, New York

Abstract

In the past, there have been many different image file formats, providing many different capabilities. However, the one thing they all shared was a very limited mechanism for encoding color. Storing a color image today forces developers to either take a "lowest-common-denominator" approach by using a single standard colorspace in all applications, or by using the capabilities of ICC color management at the loss of wide interoperability.

The JPEG 2000[1] file formats change all this with a new architecture for encoding the colorspace of an image. While the solution is not perfect, it greatly increases the number of colorspaces that can be encoded while maintaining a very high level of interoperability between applications.

This paper describes the color encoding architecture in the JPEG 2000 file format and shows how this new architecture meets the needs of tomorrows imaging applications.

Introduction

One of the chief jobs of a digital image file format is to tell an application how to interpret the samples of the image, and one of the primary aspects of this is the specification of the colorspace. Given the importance of the specification of color, and the magnitude of the problem introduced by incorrect interpretation, it is unfortunate that most digital image formats do not allow an application to precisely specify the colorspace of an image. When past file formats have allowed for a more precise specification of color, the flexibility in selecting a colorspace for encoding of the image data was often very limited.

The JP2 and JPX file formats change this situation by adopting a thorough and precise architecture for colorspace specification that still allows for a high degree of flexibility.

This architecture is built around three major requirements:

- That the colorspace of an image shall be precisely and accurately specified within the digital image file,
- That the creator of an image shall be allowed to select the colorspace of the image such that the image file best meets the needs of the target application,
- That any reader of that file shall be able to properly interpret the colorspace of the image and correctly display that image.

Obviously, these are lofty goals, and practicality must be kept in mind when designing a system based on such requirements. However, the color architecture in the JP2 file format makes great strides in resolving these requirements into a system that is practically implementable in any system (as required by the overall interoperability goals of the JP2 file format). Application requirements that are not fully met by the JP2 file format will rely on the JPX file format. This prevents many interoperability problems by guaranteeing that all JP2 files can be read by all JP2 readers, while still providing a means by which the extended features can be used.

There are two major parts of the color specification architecture in the JPEG 2000 standard: specification of color encoding and specification of color interpretation, as shown in Figure 1:

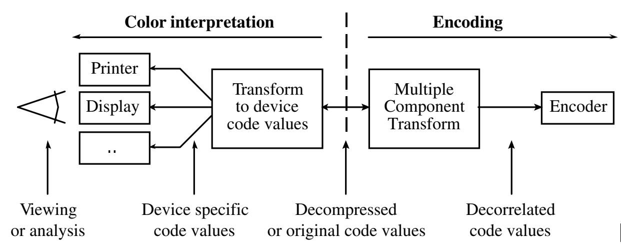

Figure 1. Division of color architecture in JPEG 2000

This architecture clearly differentiates the steps that were taken by the file writer in order to maximize compression efficiency (encoding) from those steps that must be taken by a file reader in order to properly use the image (interpretation).

Encoding the Color Data with JPEG 2000 Compression

The encoding half of the architecture represents those steps that the file writer took to prepare the image data to be compressed using the JPEG 2000 algorithm. In general, this involves a transform that mixes the components to decorrelate the data. The decorrelated data is then compressed on a component-by-component basis.

Baseline Color Encoding (JP2)

In JPEG 2000, encoding, through the multiple component transform, can be specified on a tile-by-tile basis, and thus is specified as part of the codestream syntax, not the file format. Part I of the standard defines two multiple component transforms:

The irreversible component transform (ICT) is better known as the YC_bC_r transformation from baseline DCT-based JPEG. It transforms data in an RGB colorspace to a luminance-chrominance form and is useful in lossy compression applications. Once the components have been transformed and compressed, the chrominance components can be truncated to improve compression with minimal loss of quality.

The reversible component transform (RCT) provides decorrelation of RGB data in situations where lossless compression is required. Once the components are decompressed, the reverse RCT can be applied to get back to the original code values. Obviously, if the chrominance components were truncated, there would still be loss associated with the image.

While these two transforms do not meet the needs of all applications, they are generally sufficient for digital photography and Internet imaging applications.

Extended Encoding (JPX)

The JPX file format (along with extensions to the compressed codestream as defined in Part II of the standard) extends the encoding architecture to allow for two new capabilities. The first new feature is the ability to use a custom multiple component transform. This would allow an encoder to generate an optimal multiple component transform for a particular image, thus improving quality or increasing compression. The second new feature is the ability to shape or gamma encode the image data as part of the encoding process. For example, linear encoded data is required in many high-end compositing applications; however, it compresses very poorly. The file reader would extract the inverse shaper or gamma value from the codestream and use that value to transform the image data back to its original form.

One possible image processing chain for extended encoding is shown in Figure 2:

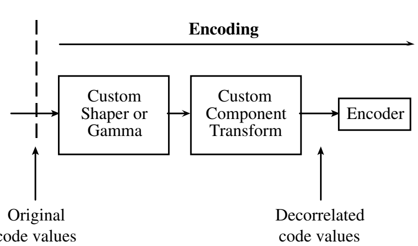

Figure 2. Encoding a color image using extended encoding options

Color Interpretation

The color interpretation half of the architecture represents those steps that a file reader must take to properly print, display, or prepare the image for analysis. In general, this involves a transform that converts the decompressed (and recorrelated) code values into a colorspace designed for use with the target output device or analysis. The JP2 and JPX file formats define several methods by which the colorspace of the image can be specified. These methods balance the required flexibility with the barriers to adoption and interoperability between readers.

The Enumerated Method

Like many other digital image file formats, the JP2 and JPX formats define a list of standard colorspaces and assigns integer codes to represent those spaces. For an image encoded in one of those standard spaces, the file would embed that integer code into the file (the enumerated method). While there are a large number of known and commonly used colorspaces, each enumerated colorspace must be natively understood by the reader, and thus the complexity of a reader implementation is somewhat proportional to the number of colorspaces that must be understood through enumeration. The image chain for the enumerated method is shown in Figure 3. Note that each input to device color transform must be natively known to the application (or how to generate it must be known).

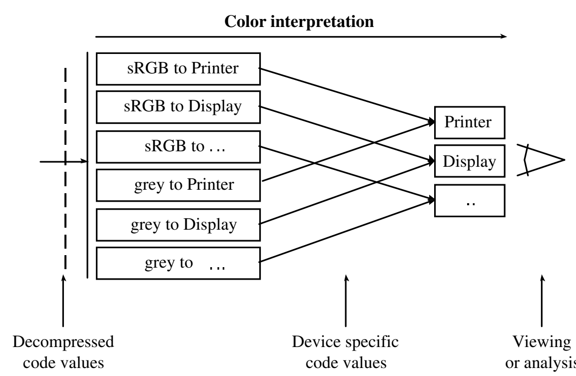

Figure 3. Interpreting the colorspace of an image using the Enumerated method

To ensure that the JP2 format is practically implementable in all applications, the set of enumerated spaces is restricted to sRGB[2] and a gray scale space related to sRGB for JP2 files. Other spaces, such as the ITU CIELAB, are expected to be standardized as extensions in the JPX file format. The JPEG committee also intends to provide a mechanism by which vendors can register and use other standard colorspaces.

The Any ICC Method

Another common way to specify the colorspace of the image is to embed an ICC profile[3] in the file (the Any ICC method). This profile (an input profile) specifies the transformation between the decompressed code values and the Profile Connection Space (PCS). To convert to a device specific colorspace, the input profile is combined with an output profile (a profile that specifies the transformation from the PCS to the device specific colorspace). The decompressed code values are then processed through the combined profile by an ICC color management engine. The image chain for the Any ICC method is shown in Figure 4.

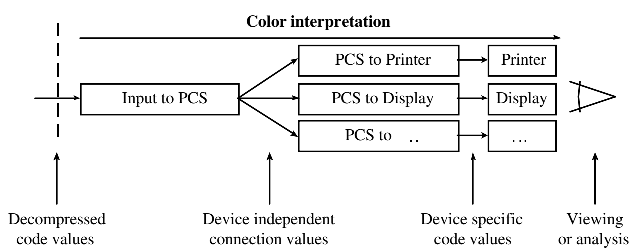

Figure 4. Interpreting the colorspace of the image using the Any ICC or Restricted ICC method

However, some ICC profiles are very complex, containing multiple 1D-LUTs, matrices and 3D-LUTs. In many applications, it is impractical or impossible to implement a fully compliant ICC color management engine. As such, the Any ICC method is expected to be defined as part of JPX in part 2, not in JP2, because it is very important that all JP2 readers be able to read all JP2 files.

The Restricted ICC Method

The major problem with the enumerated method is lack of flexibility. While sRGB is appropriate for some applications, other colorspaces are needed in other applications. For example, RGB images targeted at wide-gamut output devices should be stored in a wide-gamut colorspace, such as ROMM RGB.[4] Other images, such as those captured by scanning a consumer 35 mm negative, require the storage of information whiter than a perfect diffuse white reflector (a consumer negative may contain information up to 30,000% white). For those images, a colorspace such as ERIMM RGB[4] would be appropriate.

However, as noted above, it is impractical to enumerate a large number of spaces, as the complexity of a JP2 reader increases with the number of required enumerated colorspaces. The generic use of ICC profiles does not provide the answer either, as a complete ICC color management engine would represent barriers to adoption in many applications.

However, the ICC Profile Format Specification[3] does define two classes of profiles. These classes, Three-Component Matrix-Based Input and Monochrome Input profiles, can be implemented very easily. Those transforms contain at maximum three 1D look-up tables and a 3×3 matrix. These two profile classes allow for the specification of a very large number of RGB and gray scale colorspaces.

The Restricted ICC method for specifying the colorspace of an image in the JP2 format uses any profile conforming to one of these two classes of profiles. That profile is embedded into the JP2 file. Applications reading the JP2 file have two choices for interpreting the color of the image:

Extract the profile and use any compliant ICC color management engine to transform the image into a colorspace appropriate for the desired output device

Extract the profile, extract the look-up tables and matrix from the profile, and transform the image using application specific code. If this choice is made, it is important to note that the transformation specified within the profile transforms the decompressed code values to the Profile Connection Space. The application must combine that transformation with the appropriate transformation for generating values appropriate for the desired output device. Additional information for performing such a transformation is contained within the JPEG 2000 standard itself.

Note that the image chain for the Restricted ICC method is identical to that of the Any ICC method as shown in Figure 4. The difference is that the input to PCS transform is guaranteed to be of low complexity.

The Vendor Color Method

While the use of ICC profiles provides a very high degree of flexibility, there are two drawbacks to those methods:

- It is a non-trivial process to match a profile against a set of known profiles.
- ICC profiles are generally limited to photographic applications.

Because of these drawbacks, the JPX format is expected to define a fourth method, called the Vendor Color method, to allow any application in any domain to define a "shortcut-code" to represent a particular colorspace definition. This code is stored as a UUID[5] (Universally Unique ID); a 16-byte number that when generated correctly, should be unique "throughout the universe." This can be done without involving either the JPEG committee or another third-party.

Using Multiple Colorspace Specification Methods

It is important to note that there is often more than one way to specify the colorspace of an image. For example, a well-defined colorspace can be specified through enumeration or can be specified through an ICC profile (which technically specifies the default transformation from that colorspace to the PCS, not the colorspace itself). However, while these multiple methods often produce identical results, they are optimized for different applications. Each of the defined methods has pros and cons:

The Enumerated and Vendor Color methods allow for quick recognition, but the application must recognize the code and natively understand how to transform the image data to the relevant device colorspaces. In addition, colorspaces defined in Part II of the standard or by registration will not be understood by a conforming JP2 reader.

The Any ICC method allows almost any colorspace to be precisely and accurately specified in a generic way, but can be impractical or impossible to support in some applications.

The Restricted ICC method allows for many RGB and Monochrome spaces to be specified in a generic way that is practically implementable in all applications, but approximations may need to be made in order to represent the transformation to the PCS within the profile restrictions. However, the performance of a system (in terms of decoding speed) using Restricted ICC profiles will be higher than a system using the more complex profiles.

When creating optimal files for a particular application, the writer may be able to choose one best method for specifying the colorspace of the image. However, many files must also be interoperable outside that particular application, and performance issues may need to be considered. When considering target optimization, interoperability, and performance, a balance will often only be found by using multiple representations of the colorspace of the image.

The JP2 and JPX file formats allow for any number of different colorspace methods to be used within a single file, allowing the file writer to address these three issues, provided that all of those methods are "equivalent." In fact, it is expected that many JPX files will contain multiple methods by default.

For example, consider the registration of standard and vendor colorspaces in Part II. In most cases, it is desirable to create files that can be read by any reader (and thus meet the conformance requirements for a JP2 reader). This can be a problem, because the new colorspaces are not understood by older and simpler readers. If interoperability is indeed required, the file writer must use the Restricted ICC method to specify the colorspace of that image. However, an application that has an optimized processing path for the new colorspace would need to compare the ICC profile byte for byte with a reference profile. In the best case, this step is just a nuisance. In the worst case, such as with scene spaces such as ERIMM RGB,[4] this step is not possible, as there may not be a specific reference profile, because the profile embedded within the file must specify the desired rendering (look) of the image in addition to the colorimetric transformation.

In these cases, it is very desirable to also embed the enumerated code for that new colorspace in the file (at the insignificant cost of around 15–27 bytes).

In another example, the "best" transformation from the decompressed code values to the PCS requires the use of 3D look-up tables, which unfortunately reduces the performance of the system. In this example, the application desires to provide a "quick-and-dirty" user-selectable setting, indicating that the application should approximate the colorspace transformation to improve performance, at the expense of quality. In this example, the file would contain both the best ICC profile, using the Any ICC method, and a less complex profile, using the Restricted ICC method.

It is expected that some applications will have requirements from both examples, and thus use all three colorspace methods within the same file, as shown in the image chain in Figure 5.

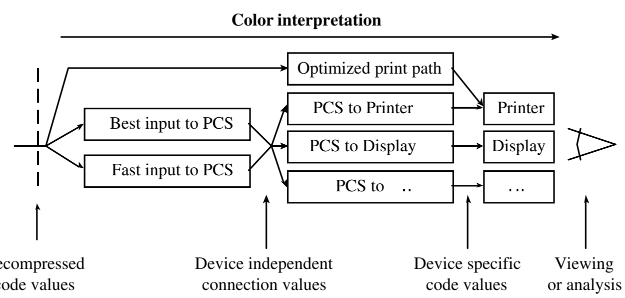

Figure 5. Color interpretation image chain with multiple options

Conclusion

The JPEG 2000 family of file formats (JP2 and JPX) provides a flexible architecture for unambiguously specifying the colorspace of an image. Using a combination of the Enumerated, Any ICC and Restricted ICC methods, an application can create files that can be both optimized for a target use, yet fully interoperable with other applications. By using the Restricted ICC method as the cornerstone of the architecture, we can ensure that all readers can properly interpret all files, even among applications that do not support ICC color management.

References

1. ISO/IEC 15444–1, Information technology—JPEG 2000 image coding system.
2. International Electrotechnical Commission. Color management in multimedia systems: Part 2: Color Management, Part 2–1: Default RGB color space—sRGB. IEC 61966–2–1 1998. 9 October 1998.
3. International Color Consortium, ICC profiles format specification. ICC.1:1998-09.
4. K. E. Spaulding, G. J. Woolfe, and E. J. Giorgianni, "Reference Input/Output Medium Metric RGB color encodings (RIMM/ROMM RGB)", *Proc. PICS 2000 Conference*, March 26-29, 2000, Portland, OR.
5. ISO/IEC 11578:1996 Information technology—Open Systems Interconnection—Remote Procedure Call.

Biography

Scott Houchin earned both B.S. and M.S. degrees in Computer Engineering from RIT, which led to a career at Eastman Kodak Company. For the last six years, Scott has been concentrated on image file formats, including working on the Photo CD system and serving as the chief technical architect of the *Flashpix* format.

In 1998, Scott was instrumental in launching a group to develop file format technology for the upcoming JPEG 2000 standard, and has served as chair of the newly formed file format subgroup within ISO/IEC JTC1/SC29/WG1 throughout the standardization process.