An Effective Colour Management Architecture for Graphic Arts

Tony Johnson*

Keywords: Colour, reproduction, printing, specification, standardisation

Abstract: This paper discusses the background to colour management and how the current ICC architecture may be improved to help meet Graphic Arts workflow and quality requirements – both in terms of cross-media reproduction and matching between proofs and prints. It suggests that an architecture that separates the various components of a single profile and concatenates the string of resultant profiles by a well-defined 'dumb' CMM would be far more versatile for cross-media reproduction, and discusses what is needed in the way of reference gamuts, profile information and control procedures to ensure colour matching. It also briefly reviews the desirability of baseline profiles that some proponents feel will be necessary to obtain consistency for some applications and what Graphic Arts will need to replace these. Finally a brief discussion is given as to how separation specific information (such as UCR/GCR) may be preserved when colour management is undertaken

Introduction

For many years the Graphic Arts industry has attempted to move the process of colour reproduction from an entirely craft based process to a manufacturing process based on fully quantifiable data. I look on the classic theoretical work of Hardy and Wurzburg (1937), Neugebauer (1937) and Yule (1938a and b) as the cornerstones of that effort and the theory developed at that time still finds application today. Hardy and Wurzburg specified the requirements for an 'ideal' set of primaries against which colour correction requirements could be defined, following which Yule developed a mathematical formulation of this procedure and showed how this could be used to define the characteristics of the colour correcting masks required. Although subsequent work showed limitations to the theory for turbid media such as printing inks, it still proved effective in enabling masking procedures to be developed that would achieve a colour match between

* Professor of Digital Imaging, London College of Printing

original and reproduction. Yule also discussed how the black printer could be defined to both expand the colour gamut and enable UCR or GCR to be achieved. Neugebauer, on the other hand, developed a different approach to achieving a colour match in which he proposed a model defining the additive mixture of the 'micro' elements of a halftone screen in terms of the recently introduced CIE system of colour measurement. This enabled the prediction of the halftone dot area required to match a colorimetrically defined colour.

Inevitably computer systems were introduced to solve these mathematical expressions and Murray and Morse (1941) produced a scanner based on solving Yule's masking equations while Hardy and Wurzburg (1948) employed the Neugebauer equations for their solution. However, the main limitation of these approaches was that they specified the relationship between the colour of the original and the print as being that one should match the other. But, in general, it is not possible to reproduce an exact copy of an image by printing, simply because the gamut of colours found in the original image cannot be matched. Furthermore, there are various issues in a cross-media comparison which mean that any colorimetric measurement of the original image should not be replicated on the print, if the two are required to look alike. The result of these issues is to produce a fairly complex colour transformation if the best approximation to the original image is to be achieved, which becomes slightly more complicated if artefacts of the original image such as colour casts and exposure errors need to be removed. The reasons for this complexity have been discussed in so many TAGA papers that it is impossible to list them all. For a short list, based only on my own contributions, see Johnson (1989, 1990, 1992a and b and 1995) as well as papers published elsewhere (1982 and 1991).

Thus the colour scanners that evolved from the early ones quoted above developed algorithms to compute, at the time of scanning, the amount of ink required to produce a specific colour when printed, even though, in general, that will not be the colour of the original image. These scanners enabled the user to empirically manipulate the ink amounts, often differently in different colour regions, in order to achieve the desired colour. Typical algorithms used for achieving this (tone reproduction, grey balance, colour correction and black printer construction) are defined in Johnson (1992a) which also discusses how these separation equations empirically circumvented the various issues which preclude colorimetric equivalence in cross-media reproduction - gamut mapping, appearance matching and measurement effects. However, these issues have received a lot of research attention in recent years and a better objective understanding is now forthcoming. It is in this context that the discussion in this paper should be seen.

But, it should not be concluded from the discussion above that colour matching is never an issue in Graphic Arts. It is a major requirement. But, it is not the match between original and reproduction that needs to be exact but between the printed copies and proofs made in various sites. This requirement led to an almost universal workflow in which the creator of the separations interpreted the original as required by the customer and provided a proof of this interpretation for his approval. The matching requirement came when advertisers provided copies of a set of separations to a number of printers for them to replicate the proof. This led to the development of 'standard' conditions for printing and/or proofing so that the results obtained from such separations would be consistent between the various printing sites and match the proof. So long as each of the printers could match the approved proof consistency between them was assured. All that was then necessary was that each printer maintain this result during the course of the production run. During the past 30 years these 'standards' have evolved for various processes around the world, and during the past decade a number of these have been consolidated into a series of International Standards.

For such a situation to be fully effective it is necessary to ensure that proofing systems are predictable and provide a good simulation of the print. Much research and development effort was put into the problems of matching flat bed proofing presses to the running press and vendors of 'traditional' off-press proofing systems sold their product as a good match to the press. However, since these proofing systems were simply reproducing the separations it was essential that the colour of the pigments (or dyes) matched those of the inks and that the tone reproduction of the proof simulated that obtained by platemaking and printing.

More recently, however, the industry has progressed towards the transportation of electronic images, rather than images being provided on film. Such a situation permits other workflow scenarios to arise. Currently, image exchange is still largely based on the assumption that the files moved between users are encoded in CMYK format, prepared for a specific printing process; such as one of the standard processes described above. The files could be thought of simply as electronic separations; although generally in continuous tone form, without the imposition of the halftone screen. Thus, all the constraints placed on controlling the process are the same as for the exchange of film separations. It is obviously essential for the final user to ensure that the output device is properly calibrated, and that the correct halftone screen is selected, but otherwise there is little difference to the workflow based on separations.

However, the exchange of electronic files also offers new possibilities - the ability to process electronic files with a colour transformation that can be implemented at any stage in the process. This offers a variety of new possibilities, and dangers. The advantage of this approach is the possibility of reproducing a single encoding of an image across different output devices with quite different imaging characteristics (substrates, colorants and screening techniques). This leads to two possibilities – repurposing the data for various applications other than traditional Graphic Arts (which in general will not produce a colour match between the reproductions because of gamut differences) and the possibility of reproducing a matching copy of the image across a range of devices, where the gamut is well defined. From a printing and proofing standpoint this latter possibility means that it is far less important that the process characterisation be consistent between printing presses as it was when film separations were sent to each printer. Each individual press can be 'fingerprinted' with its own characterisation data, such that a colour transformation may be applied which matches the proof, and yet is unique to that press, when the image is produced. The reference is no longer the halftone dots themselves but the colour that they will produce. This allows greater flexibility in the choice of media and printing conditions, which may lead to improved consistency and production. There are only two requirements for the colour match to be achieved – the ability to define the transformation and the capability of reproducing the gamut of colours produced on the proof. The control procedure for this approach is usually known as colour management and it is usually assumed – though not necessary – that the transformation will be based on colorimetry.

Thus, colour management can be differentiated from the traditional approach, which is still used by many requiring the highest levels of quality, by the reference colour space used. For the traditional approach the colour space, often defined by density values, was specified differently by each equipment vendor. Based on this, as discussed earlier, a series of empirical equations were defined to enable separation and the operator was provided with controls to change the parameters for the specific printing conditions encountered. Colour management, on the other hand, usually attempts to use the CIE system of colorimetry as the basis of its colour transformations. The relationship between this and each of the scanner RGB values and the printed CMYK values are defined, typically by look up tables derived from colour measurement of known combinations of device values. These separate tables can then be combined (as they have a common CIE reference) to produce the overall transformation that enables separation. In order for this to be successful the tables need to include all information necessary to achieve appearance matching (including measurement artefacts across different media types) and gamut mapping.

From the discussion above we can conclude that Graphic Arts has three distinct needs from colour management. These are:

- Defining general cross-media reproduction (where the gamuts and viewing conditions are often not the same and therefore require solution to the problems of appearance modelling, gamut mapping and any measurement artefacts).
- Exchange of images to multiple sites which all need to match for colour when reproduced (which either requires exchange of the data to occur in an encoding space in which no colour ambiguity is possible - or agreement on how all the parameters used to define the reproduction are exchanged).

- Maintenance of separation specific information that cannot be defined colorimetrically (such as the GCR level or the requirement for additional colorants) where maintaining this is deemed important.

At the present time the majority of users of colour management in Graphic Arts seem to concentrate on applying it to the matching of proofs and prints, and not to the general cross-media reproduction situation. This is because that requirement is relatively easy to define unambiguously and has thereby enabled relatively low cost devices to be used for proofing. However, there is undoubtedly some desire to see it used in the more general cross-media environment and so much of the following discussion will concentrate on what needs to be done to make this more useful to Graphic Arts before returning to the other issues listed above.

Colour Management Architecture

Most colour management currently being undertaken relies on specifications defined by the International Color Consortium (ICC). The basic architecture currently defined by ICC is really quite simple. At the centre is a reference colour space known as the Profile Connection Space (PCS). It is defined such that images are specified by their CIE tristimulus values in relation to the typical graphic arts viewing and measuring conditions, specified in ISO/FDIS 3664 (2000) and ISO 13655 (1996), for reflection copy with an unlimited colour gamut. This means that the image pixels can be encoded directly by their tristimulus values (or CIELAB values) - specified for the conditions; D50 illuminant, 0°/45° measurement geometry and 2° standard colorimetric observer. The use of a media with an unlimited gamut means that all realisable colours can be specified quite easily and the compression to any real gamut is defined at the time the image is output. However, there are potential quantisation problems by such encoding. So, the most important aspect of the ICC specification is to define a mechanism whereby images may be encoded as device dependent data, but with provision of a look-up table (profile) which is tagged to the image and defines how this data may be converted into such tristimulus values, as is the more common practice.

So, if the image is rendered such that the tristimulus values describing each pixel are achieved on output it will reproduce the colorimetric values of the original image exactly. However, it may not then look correct because of the appearance, gamut and measurement issues mentioned earlier. To, accommodate this the specification states that if the image is to be editorially corrected (e.g. a cast removed or a tonal adjustment made), or compensation made for media and viewing condition differences, this should be incorporated into the input profile which defines how the scanned data should be transformed to provide the correct values in the PCS. In this way the PCS values specified define the tristimulus values necessary to provide the desired **appearance** of the image to be reproduced on the reference print media. Thus, the highly objective colorimetric definition is replaced by a more subjective characterisation, which corrects for appearance and measurement effects, but has certain disadvantages that I will return to later.

The basic ICC architecture is shown in figure 1. The X, Y and Z values (which may also be CIELAB values) that provide the common reference for input and output are defined according to the requirements of the PCS. As already stated, particular measurement conditions are specified to ensure that the tristimulus values have a properly defined reference. The architecture defined by the specification assumes that the input profile incorporates device characterisation, as well as any corrections for measurement and appearance effects. It may also include any editing corrections, if desired. The output profile incorporates both device characterisation and gamut compression, as well as correcting for any measurement and appearance effects. Thus, the ICC requirement for colour transformation, when rendering images to any output device, is that by simply combining the two profiles into a single transformation, which is then applied to the image, should ensure it is properly reproduced.

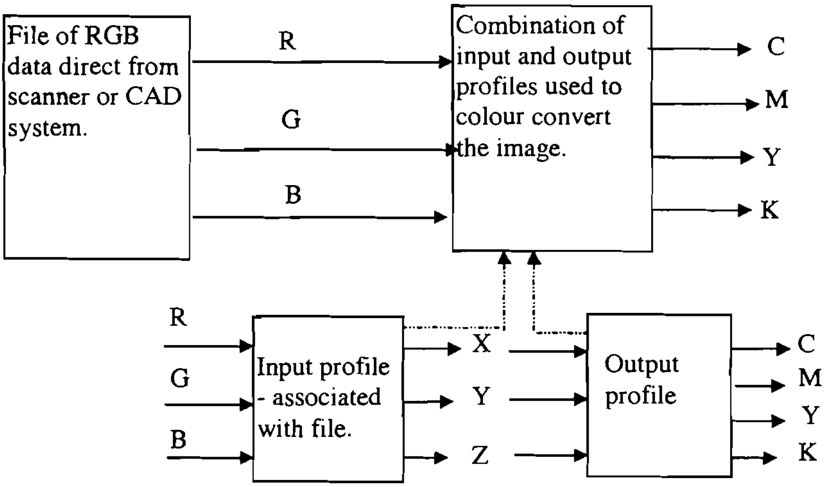

Figure 1 – Schematic diagram showing a typical ICC colour management procedure

The description above describes the general architecture, which is primarily directed at the cross-media reproduction situation. It assumes that the appearance, gamut mapping and measurement issues referred to earlier can be defined in such a way that they can be incorporated into profiles – with the current ICC architecture assuming that the gamut mapping is specified in the output profile. A number of gamut mappings are specified by ICC, of which the perceptual and colorimetric renderings are most appropriate to Graphic Arts. Perceptual rendering is that appropriate to the general cross-media rendering –

when the gamuts are inevitably different – but, also inevitably, that is not defined in any way. Thus, the results obtained will depend significantly on the system used for profile generation.

Colorimetric rendering, which is appropriate to the colour matching requirement discussed earlier, also requires the combining of profiles to convert image data. However, the profiles used should be those that produce a colorimetric match and thus the gamut issue should not arise (if the overall system is well defined) and neither will the appearance issue, in general. A more detailed discussion of the issues to achieve this will be returned to later. In principle such a transformation should be unambiguous although different methods for building profiles, and differences between the CMMs that combine them, mean that this is not always correct. However, as time goes by this situation will improve.

By providing a reference PCS one of two options was selected by ICC. The alternative would have been to simply provide colorimetric data within the profile for whatever conditions of measurement and viewing were used for the media – together with a specification of the conditions for which the colorimetric data was computed. However, that has certain difficulties in transforming data to a common colorimetric domain if different measurement characteristics have been used for characterising devices. Whist there are reasonably good models for correction for chromatic adaptation (defined by the colour of the source used for viewing) it is less easy to correct for different measurement geometries and the other standard observer. Clearly by defining a PCS, and the associated measurement conditions, that difficulty is overcome.

However, the disadvantage of such a specific PCS is that for a user who has very different, and somewhat limited requirements, such a transformation produces an unnecessary complication. Whenever, the output media is not a print but, say, a monitor unnecessary chromatic adaptation corrections will have to be made for both input and output. This is why using such a PCS can be problematic when no print is involved in the reproduction process (or it is not to be viewed under the conditions specified in the PCS) since subjective corrections are made to the characterisation, which are not separately specified (in the current architecture). Thus, we can conclude that characterisations based on some clearly defined PCS can provide an approach that is useful for the general cross-media situation - where prints are involved that are viewed under D50, at 2000 lux - but provides unnecessary complexity for other requirements.

If a PCS is necessary it is reasonable to ask whether the specification selected is the best choice available. For Graphic Arts it is probably as good a choice as any. The viewing illuminant and measurement condition selected are those most directly applicable to Graphic Arts and the only real issue from my perspective is the choice of a print medium with unlimited colour gamut as the reference. There is also something of a problem in that the 2000 lux illumination level is very high and the current version of ISO 3664 provides a 500 lux level, which is more practical for the assessment of images. However, since the 2000 lux level will also be used, for comparing prints and proofs, both levels need to be accommodated so any selection can only be arbitrary and 2000 lux is as good a choice as 500!

In my view the reference print medium selected for the PCS makes profile evaluation quite difficult, but more importantly it leads to a great deal of ambiguity. The issue is what values need to be encoded in the PCS for the white and black points of any real image, particularly due to uncertainties around the issue of viewing flare and the need to allow for catchlights in an image. Simply defining a viewing condition and measurement condition – without attempting to relate the two – does not properly define the appearance of the image. By defining a real medium – that can be visually assessed under the specified viewing condition – allows that problem to be overcome. Thus, I believe that if a PCS is retained in any colour management architecture a realisable, large gamut, medium should be specified.

My major concern with the architecture as defined above is that it confuses too many issues within a single profile. I believe there are distinct advantages in keeping the various components of the profile separated. Thus, if the PCS model approach is followed, an input profile should be composed of the following sub-profiles which are defined quite separately:

characterisation data for the device and media,
an appearance correction for the media viewing conditions,
any measurement correction to accommodate viewing flare for the media,
a gamut mapping (if for no other reason than the need to 'scale' the dynamic range to properly place the image white and black onto those of the reference connection space), and
any global image editing information.

A similar set of transformations would be provided in the output profile although without the editing transform, which seems superfluous at this stage.

I would anticipate that the specification would be defined such that the default condition would be that the CMM combines these components without reference to the user. In that sense the behaviour to the user would be simple unless specific changes are invoked. Also, it should be clear from this description that I believe it is desirable to define both image and media transformations in the profile.

The reason for the inclusion of both image and media transforms should be reasonably clear. The media transforms are essentially fixed for a given media and ensure colorimetric or appearance accuracy (whichever is required) for every image. However, quite often neither is appropriate. An extreme example would be the reproduction of colour negatives where a colorimetric or appearance match would be somewhat bizarre. However, a somewhat more subtle (though equally important) deviation is when images (particularly colour transparencies) are improperly exposed and/or have a colour cast. Reproducing these images to maintain these artefacts would not normally be ideal.

One way of overcoming this is to have image global editing information – such as gamut mapping information (particularly white and black point placements) and colour cast correction - specified within the profile. It is clearly not **essential** that it be part of the input profile since it can always be calculated from the image later – thereby leaving it to the receiver to do. But to avoid ambiguity across multiple devices it is best that instructions be given to them directly. Even then it may not need to be incorporated in the profile – with the user-interfaces provided on some scanners it is often much simpler to determine at the time of scanning in order to 'correct' the image data itself. Correcting the data on the scanner, where higher precision may be provided, ensures that the most efficient quantisation is achieved. However, where no such facility is provided on the scanner incorporating such corrections in the profile makes sense from a quantisation standpoint. With the current architecture it is necessary to make these corrections to the image data itself – or that they be included in the profile, but not separately defined. Thus, I feel that in this respect the current architecture is sensible – all I am proposing that is different is to make the individual transforms explicit.

However, a somewhat simpler model than that proposed above is achieved if the PCS concept is no longer mandated. Essentially the conversion from input device space to output device space requires an editing specification, a pair of colorimetric definitions (but not necessarily to a common reference), a pair of colour appearance definitions (to accommodate viewing conditions and measurement effects) and a gamut mapping. The way that this would work is shown in figure 2. This is essentially the 5 stage transform discussed in Johnson (1992a), but appended with an editing profile for the reasons discussed above. The only issues this raises are those described earlier – the problems of correcting for different measurement geometries and observer. Ideally such corrections would be included in the appearance definitions but at present I am not aware that any general transform exists. Thus, for the time being it may be necessary to mandate a specific measurement geometry, sample backing and observer.

The basic parameters used in the colorimetric and appearance specifications should be defined in the tags accompanying the data. In this way the receiver should be told what assumptions have been made. However, it should not preclude editing the results of these calculations where the models used have not provided the necessary quality. Thus the receiver needs to know what appearance and colorimetric assumptions were made so that it is easy to decide

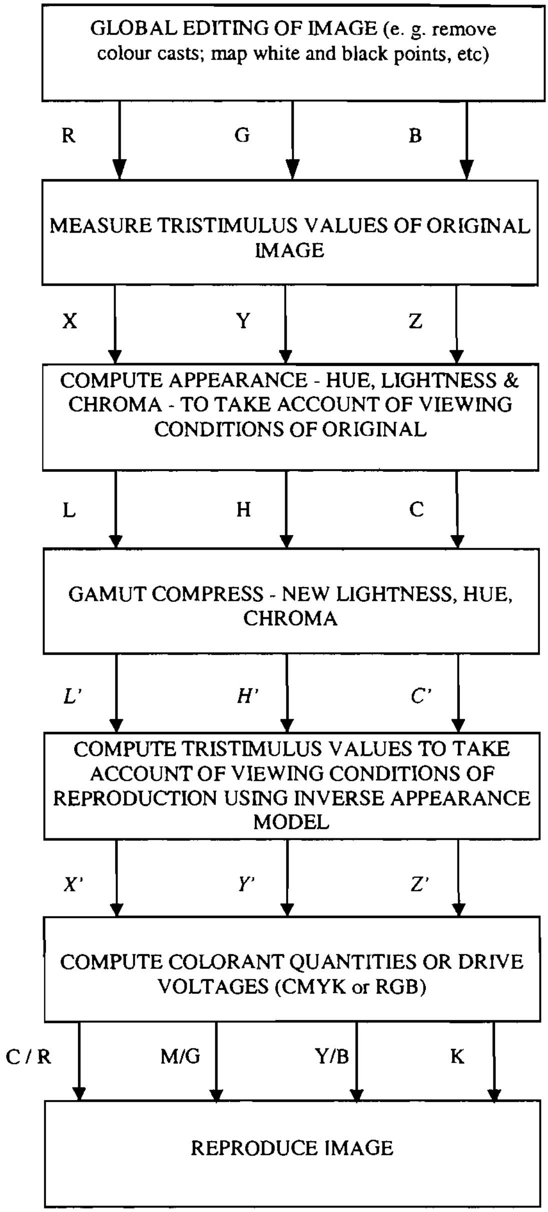

Figure 2 – Schematic diagram showing the calculation procedure without a PCS

when, and how, to replace these. But it is not necessary that the results of the look-up table produce exactly the same information – this would be far too restrictive from a quality standpoint.

It should be noted that in figure 2 the image gamut mapping and media gamut mapping are separate. The former occurs in the first stage and the latter in the fourth stage. Based on my own research, Johnson (1977), this is adequate providing the white and black of the image are properly selected and the neutral scale is linearly scaled in 'appearance corrected Lightness' for both image and media gamut mapping. Thus, with such an architecture, stages 2 to 6 remain fixed for a specific device pair and only the first stage varies according to the image. However, other researchers claim that different lightness renderings are required that depend on image content (e.g. Jorgensen (1976), Archer (1985) and Braun and Fairchild (1999)). This could be interpreted as requiring that the nature of the overall gamut mapping is both image and media dependent simultaneously and needs to differ when, for example, subjects are high-key or low-key. In such cases the gamut mapping in stage 4 would also need to be based on both image and media.

There are a number of advantages to the approach shown in figure 2:

- it will enable any user who has rendered an image on one device to provide the editing, appearance and gamut mapping information necessary to ensure it is matched by another that has quite different characteristics (even if it has the same gamut). And yet, at the same time, if it is necessary to alter any of these parameters (e.g. to accommodate a change in viewing condition) this approach provides the flexibility to do so. Combining appearance, editing, gamut mapping and characterisation, as is currently done in the ICC profile, means that a whole new profile has to be built to alter any of these since they cannot be separated and the receiver, who is most likely to do this, will find it very difficult to replicate what was done by the sender.
- it is really only possible to define a gamut mapping when both the input and output gamut are known (and possibly the image). By combining gamut mapping with the output profile, separated by the PCS, it is clearly not possible to do this. It will be fixed regardless of the input gamut and, in my experience, this generally means that profiles need some editing whenever specific pairs of profiles are encountered. Thus users still require device-pair specific mappings, for the highest levels of quality, which have to be treated as a pair. I believe that a well-defined concatenation of sub-profiles makes it easier to identify the specific causes of any problems and optimise those directly – thereby reducing the number of device specific profile pairs required.
- the inclusion of an edits profile ensures quantisation artefacts are minimised for many scanners.

One of the problems I see with the architecture defined above is that it is potentially more complex to use. How it is used depends upon the applications. For some applications consistency across different media is important, for others it is desirable to obtain the optimum reproduction on each media. Sometimes simplicity is important and other times the added value to the product is such that users are willing to expend considerable effort in optimising their system. This raises the issue as to how any architecture can meet all sets of requirements. To achieve this it is important that software developers ensure that the complexity is minimised by providing 'default' modes in which transforms are based on previously defined 'sub-profiles', and that good user-interfaces are provided for editing 'sub-profiles'. At the same time Graphic Arts developers will need to provide well designed analytical tools for optimising the system.

Of course, in order to make such a procedure work effectively it is essential that tools be developed that enable users to build and edit the various sub-profiles described earlier. My expectation is that the sub-profile format should be specified to be consistent with the existing profile specification so far as possible – so that existing profile editing tools can be modified relatively easily to edit them. It may be advantageous for developers to offer tools that modify appearance and gamut mapping algorithms parametrically, although I think that approach will make them too complex for typical users.

The approach proposed by some ICC developers who have been exploring ways of improving the ICC architecture is to have a 'smart CMM' that does more than simply combine profiles – it actually implements gamut mapping, appearance modelling, etc algorithmically. The intent would be, presumably, to base these on the 'best' knowledge currently available. Thus the use of CIECAM97s for appearance modelling and one of the better gamut mapping algorithms (if one is not specified by CIE 8-03 in the near future) could be identified. This obviously makes sense for the cross-media reproduction requirement and will undoubtedly meet the needs of many users. However, such default (or baseline) procedures are unlikely to meet the needs of Graphic Arts and the system must allow alternate algorithms to be used. But, that 'smart CMM' approach will become a problem when matching is required so it must also be possible to define the result of these algorithms as tables that can be moved with an image, and that the CMM can still simply interpolate and combine tables that define non-default versions of each of the procedures specified earlier.

Making colour management work for Graphic Arts

The discussion above is focussed on the cross-media issues and that is where most of the complexity and architectural issues arise. However, as stated earlier exact matching of images across different devices (possibly with different media) is also of fundamental importance to Graphic Arts. The reasons for this centre around the importance of the proof in any sensible workflow, and the subsequent need to be able to print identical images in multiple sites. The purpose of a proof print is to simulate the visual characteristics of the finished print product as closely as possible and is therefore an essential element to Graphic Arts colour reproduction. It should show the creator of the image supplied for reproduction how it will actually look when taking account of edits required on the original (such as cast removal, local colour changes and contrast enhancements) and the limited gamut which is likely to be used for rendering. Except for the simplest of situations it is largely impossible for the user to define all these changes in a simple numerical way and some degree of subjective interpretation is essential. The only way that the print buyer can actually see this, even if he scans and edits the image himself, is to view a proof of that image to determine whether the result is pleasing for the purpose. It is this requirement that places certain, quite restrictive, conditions on the whole process.

The historical colour reproduction approach has been to ensure that the colour separator, proofer and printer have previously agreed to control their processes according to a set of parameters that uniquely define the visual characteristics of both print and proof. Such an agreement enables the production of suitable separation data (without recourse to "trial-and-error") and subsequent production of proofs and prints from this separation data. The model used has been to assume the transfer of separation films between sites, as discussed earlier, and the control procedures to ensure consistent image transfer have been specified by International Standards (ISO12647 parts 1 to 6).

However, unlike those standards, a colour management approach assumes that only electronic data will be transferred and, as discussed earlier, this may be altered according to the characteristics of the printing colorants used. The data that is exchanged may be defined as colour separated input tone values (which may be interpreted as halftones values for halftone printing and proofing processes) and currently this is by far the most common approach as explicitly specified in ISO 12639 (TIFF/IT) and CGATS draft PDF/X. As will be discussed later only colorimetric rendering profiles are normally required when using colour management in such a system, to enable the conversion of the data from the process it was specified for to any other, with the sole purpose of obtaining a colour match. In such cases it may only be necessary to know the process for which the data was specified by name (particularly if a standard process such as those defined by ISO 12647 was used) as a profile based on the standard data sets such as ANSI CGATS TR-001 (1995) is likely to be widely available.

But need we be so 'print-centric'? There are reasons for believing that images will be increasingly exchanged as unseparated data. While this offers minimal (if any) direct advantages to traditional Graphic Arts users it is clearly sensible when the image data has multiple applications. If exchanged in this format it **must** be appended with supplementary information defining the colour it is intended to represent (rather than the colour of the original image). The final images may be reproduced by digital proofing and printing techniques or by a mixture of traditional printing processes. Similarity of colour is expected across all of these, as stated above.

It follows from this that different workflow scenarios need to be considered. In one such workflow the digital files may be proofed and printed on the same device, and only on that device. In that case an agreement may be reached between originator and printer such that little may be gained from any form of standardisation. So long as the process is controlled to ensure consistency over time satisfactory results will be achieved. However, another scenario (which is the one of interest in this paper) may involve a single file (which may be print specific CMYK – but not necessarily that to be used for production – or unseparated data) being reproduced across multiple devices, utilising various marking technologies and possibly in quite different places and times. One of these reproductions may be the proof. In order to obtain a visual match between the various prints, each of the processes may require separation values (solid tone coloration and tone value increase) which are different from each other. This is caused by differences in reproduction technology between the various digital and traditional printing processes which may well introduce varying halftone or continuous tone rendering processes as well as the effects of phenomena such as gloss, light scatter (within the print substrate or the colorant), metamerism and transparency. In all cases it is important that the appropriate corrections are specified to obtain a colour match.

At face value it might appear that the use of colorimetry to define a colour match would deal with this situation (providing the gloss, scatter and transparency effects do not cause measurement issues). The proofing system would show the effect of any edits required on the image and the printing system would then show the best approximation to this proof that could be produced on it. However, unless the colour gamut of the proof is constrained to be within that of the print the printer is placed in a very difficult position. Does he try and extend the gamut for that press run, and if he cannot what is the best approximation? Bearing in mind that he is likely to be mixing images from various sources on the press forme it can be seen that it quickly becomes impossible to properly predict the final result in such a situation. I believe it is really quite important that the buyer of the print sees a proof that is limited in gamut such that it provides a reasonable approximation of the final print product. To this end the gamut of colours which can be produced on the print must be known, and taken account of, at the time of proofing.

So, in order to achieve such a match it is necessary that all the data needed to define the rendered colour is exchanged with the image file and that the rendering processes meet 3 criteria:

a satisfactory colour transformation is defined for each process that enables

- the supplied file to be rendered accurately
- that each of the printing (or proofing) processes is adequately controlled to maintain the integrity of the transformation
- that the colour gamuts of each process are similar (or restricted to an agreed gamut)

There are two implications of this discussion. The first is that the various printing gamuts need to be pre-defined in order that proofing systems can match them. The second is that the proof profile (or the appropriate gamut mapping, appearance mapping, etc if the revised architecture is considered) **must** be transmitted to the printer in order that he can achieve any specific characteristics associated with that system. To achieve this I believe that the Graphic Arts industries need to specify a number of reference printing gamuts, and also define exactly what supplementary information needs to be transmitted with any image exchanged. ICC needs to ensure that the mechanism for inclusion of this information is agreed.

The definition of a printing gamut follows directly from any standard ink and printing definitions – such as those provided in the various parts of ISO 2846 and ISO 12647 – and there seems no logical reason not to use those as the basis of any set of standard gamuts. However, it should be clear that while the paper and ink colour attributes are clearly very significant in the definition of the gamut limits the tone value increase parameters which play a large role in the ISO 12647 standards are of far lower importance in the colour management environment. The way in which the data inside of a gamut is encoded must be defined within the proof profile (which has to be transmitted with the image) but, in principle, any encoding can be accommodated. Perhaps the only workflow scenario where tone values (with tolerances) need to be considered directly is when digital halftone proofing systems are used. By definition a halftone proofing system must have colorants similar to those of the print and specifying tone value increase is a very convenient way of quantifying the way in which the halftone elements are transferred in the process.

Some proponents of the reference colour gamut approach prefer to also have a reference encoding specified which describes how the gamut is filled. McDowell (2000), for example, suggests that such an approach reduces the proliferation of profiles which will result as users each characterise their own devices and 'tune' profiles. I must say that I am not fully convinced it is **necessary** to do this – although it cannot be harmful I am not sure it is beneficial, so long as the profile actually used is appended to the image and procedures to utilise it are implemented by the receiver. McDowell argues that receivers only need to set up one profile combination with the reference encoding approach (which simplifies implementation) but I am as yet unconvinced. However, it does avoid the requirement on proofing and colour management vendors to develop new procedures for restricting gamuts beyond that of the device itself when defining a profile with perceptual rendering. Most current procedures to do this require the encoding within the gamut to be known also. Of course, if the new architecture proposed earlier is adopted such a procedure becomes far simpler.

ISO TC 130 is currently discussing precisely how many standard gamuts are required. The gamuts being considered include:

Category 1	Extra Large - Utilising more than 4 process colours and printed on a coated paper of 70 gm per square metre or greater with a brightness over 75.
Category 2	Very Large - Coated paper of 70 gm per square metre or greater with a brightness over 75.
Category 3	Large - Coated paper of 70 gm per square metre or greater with a brightness over 75, or lighter weight gloss coated paper
Category 4	Enhanced - Super calendered paper or high quality matt coated paper
Category 5	Medium - Matt coated paper or high quality uncoated paper
Category 6	Small - Uncoated paper or Improved newsprint
Category 7	Very small - Newsprint

It will be noted that the seven classes of printing identified in the list above are largely based on the paper type used, although they are also a function of the colorants used for printing. This is because with traditional printing processes the gamut achieved is largely a direct function of the substrate used - unless the print is varnished or laminated or very thick ink films are applied. This is because the gloss of the substrate has a significant on the gloss of the ink film – which in turn has a major impact on the way the specularly reflected light desaturates the colour. However, with some ink drying methods, and also for some of the non-traditional processes, the 'ink' undergoes little absorption into the substrate and the printed gloss levels tend to be less affected by the substrate. So, within the TC 130 committee discussing this there are some who feel it is better not to specify the gamuts by any reference to the substrate.

Other committee members feel that 7 classes is too many and that only 4 colour processes should be included in the range of gamuts specified. It therefore seems feasible that the final list recommended by TC130 could be as low as 5. The choice and definition of reference colour gamuts is discussed in these proceedings by Schlaepfer (2000).

The issue of the transmission of a proof profile is a slightly complicated one. If the industry adopts a colour management approach based on the current ICC architecture every image will need to have associated with it the input profile and output profile used for proofing (with both its perceptual and colorimetric renderings). At the time it is output for printing the input profile, forward proof output profile (with perceptual rendering to provide the gamut mapping), inverse proof output profile (with colorimetric rendering to generate tristimulus values) and print output profile (with colorimetric rendering) would need to be combined to ensure print and proof match. Providing the proof has been constrained to the gamut which the printer can achieve, and the CMMs used by the various printers and proofer are consistent, a colour match should be achieved. Of course it is assumed that measurement accuracy and profile generation tools are such that characterisations can be produced reasonably unambiguously – but this is a necessary pre-requisite of colour management. Unfortunately, this is not always the case. However, profile editing does enable the correction of profiles where this is an issue.

However, as already stated it would be helpful to Graphic Arts to have a colour management architecture that consists of a greater number of sub-profiles within each of the input and output profiles:

- the characterisation data
- any global editing data 'edit profile' for edits to the image
- an appearance transformation (which may be separated into appearance and measurement effects associated with the media), and
- a gamut mapping profile transformation

If this is implemented an alternative (and somewhat simpler) approach is possible when the output gamut is known at the time of scanning, as it largely is for Graphic Arts for proofing purposes. In this situation all that needs defining is the input profile and gamut mapping/appearance transformation. The receiver would then simply be required to use a colorimetric rendering on output. This is why keeping the components of the profiles separate can sometimes be more effective; it provides a far more general solution. Of course, this approach cannot be used if it is necessary to maintain the characteristics of the black printer specified at the time of proofing. However, although this may be beneficial where halftone proofing systems are used it cannot be required otherwise. If it is required I can see little advantage in not supplying pre-separated data! In fact the flexibility of leaving the UCR/GCR decision to the printer has to be a significant advantage. However, there are Graphic Arts workflows where the UCR/GCR is defined either directly (by sending separated data), or indirectly in the output profile. For such circumstances it is necessary for Graphic Arts to have black preservation algorithms for situations where the printer wishes to use colour management to transform the data for his own printing conditions but not lose the UCR/GCR content. These will be discussed later.

Yet another approach, which would not be very different from traditional work practices, would be to consider workflows with a fixed output profile for each printing condition. The implication of this is that all users in a particular market sector, printers and proofers, would use the same profile for output, which could even be standardised. The generator of the input image would then edit the input (or gamut mapping profile) to overcome any shortcomings they perceive with this output profile, and edit the image (or profile) for any image specific corrections. The proof would be the validation that this has been done properly. There are certain attractions of simplicity to this approach – if all printers wish to maintain standardised conditions of printing as they do now. But, frankly, one may just as well send separated data and avoid the same transformation taking place many times. I have heard it said that it is better to maintain the raw data (or even CIELAB data) and this is the justification for this approach - but that fails to take account of the issues of precision and quantisation. The encoding is almost immaterial in this scenario - when the output is defined precisely and uniquely it makes sense to use that for the exchange encoding space. Of course, if the fixed output profile is not printer specific – but effectively only specifies the gamut mapping - then it is no different to the recommended workflow discussed earlier

Whichever of the above workflows is adopted by the industry it is clear that some degree of standardisation of the printing conditions is important to ensuring satisfactory reproduction in a colour management environment; either simply to define the gamut limits or, as in the latter case, to define the whole profile. And, of course, it is important that the printer has defined parameters other than solid coloration (tone value increase and plate control standards) which need to be well controlled in order to maintain the characterisation he has defined.

Black (additional colour) Preservation

One of the reasons for exchanging CMYK separated data files (or unseparated data that has a specific separation profile associated with it) is to define and maintain the black printer information. Clearly it is necessary to maintain any black monochrome information (such as text, borders, line art and tints) such that it still only prints in black when transformed to another domain. However, some people suggest that it is also desirable to define the UCR/GCR levels for colour images at the time of proofing and maintain these for subsequent printing.

If only the monochrome information is considered most page descriptions enable this to be defined and maintained as black when output, regardless of the image colour encoding. However, clearly this is not possible for the UCR/GCR levels. Since UCR/GCR is primarily a printing requirement it seems illogical to me to define it at the outset – and maintain it – in situations where colour management will allow the printer to set his own preferred level of UCR/GCR. But many workflows still seem to demand this. Of course, where files are separated even the monochrome information is generally only defined by the fact that it only prints from the black plate – there is no page description for it. If a printer provided with CMYK data wants to make use of the flexibility provided by colour management then it is important that at least the monochrome black be maintained.

In such situations it is desirable to define a procedure whereby the black (or any specific colour) can be maintained, and the cyan, magenta and yellow modified accordingly. Although such algorithms were developed many years ago I am not aware that they have been written up anywhere so, although they are relatively simple, it may be useful to briefly document typical approaches here. (Those requiring more background information on black printer models may wish to refer to Johnson (1985)).

The approaches described all require that the black channel information is read into a separate file and corrected for any differences in tone value increase between the process for which it was made and the process for which it is now required. The next stage is to determine the CIELAB values of the various CMYK pixel combinations. The inverse colorimetric output profile supplied with the image will achieve this. In order to produce the conversion from CIELAB to the new CMYK a profile construction tool is required that computes an output profile with the specified black – rather than one that builds a profile using proprietary UCR/GCR algorithms.

One possible approach would be to 'standardise' how to achieve and describe UCR or GCR, and therefore have no proprietary algorithms. It would then be possible to simply specify the UCR/GCR 'amount' in the header of an image, which could then be used in the re-separation. However, as discussed by Johnson (1988), it is by no means clear that there is any advantage in such standardisation as it would be difficult to get agreement on how to describe it. Thus, since no such standardisation is likely to be agreed proprietary algorithms would not describe the black in the same way and so a simple 'descriptive' approach is not one I believe is likely to be successful.

In designing a black preservation algorithm I would note that it is particularly important to ensure that all monochrome images MUST be maintained as monochrome, but that it is less important that the UCR/GCR is reconstituted exactly. So long as the Total Dot Area is not significantly altered (which is unlikely unless massive tone value changes in the dark tones are encountered) the combinations of Black to CMY at lighter tones is less critical. Thus, it is my contention that so long as algorithms maintain this approximately that is sufficient. It is easy to define the special case that absolutely ensures the output remains monochrome, when the other values are zero, so the problem to be resolved is to approximate the UCR/GCR level.

Typical approaches that I have investigated for this are:

1. To estimate the UCR/GCR in an image, in terms of the model used by the profile building tool of the colour management system used for building the profile, by analysing a sub-sampling of pixels in the image and finding the 'best' UCR/GCR for that model by linear programming. This works well –

but only gives a good approximation to the original UCR/GCR if a lot of points are defined.

2. The Neugebauer equations as modified by Yule and Colt (1951), but extended to 16 terms to include black as described by Yule (1967), provide a satisfactory model. Since the black tone value is known the solution to the equations is straightforward – albeit iterative – and similar to that for the basic equations as described by Pobboravsky and Pearson (1972). However, as the approach normally used is to populate a 4 dimensional look-up table, which is used when processing images, the fact that it is iterative is not a significant problem. The method does provide UCR/GCR that is close to that of the original file – within the accuracy of the model.
3. Using a model based on addition of colorimetric densities, but corrected for additivity failure as described in Johnson (1985). This method also provides UCR/GCR that is close to that of the original file – again within the accuracy of the model.

Control elements and how to use them

For control of colour reproduction across devices in a colour management environment the traditional approach, based only on control of the colorants (colour of the solids and tone value increase), cannot be satisfactory with proofing and printing systems based on various colorants. Such procedures certainly provide a method for control of the individual devices themselves and usually need to be implemented as calibration procedures by the users of the devices.

However, given the relative complexity of the system of reproduction, and because it is necessary to ensure that the tolerances which must be provided for the various devices do not combine in such a way that they produce relatively large differences between the same colour reproduced on these various devices, a supplementary control procedure is required. As proposed at TAGA three years ago (Johnson (1997)) it should be possible to achieve this by a procedure in which a number of colours are specified and tolerances are placed around them. By outputting these colours on the various devices, with every image output, it will be possible to define the quality of the output profile in matching the various prints.

The approach I recommend is to define 21 colours by their position in the reference colour gamut as shown in figures 3 and 4. This defines 3 neutral colours plus a selection of colours that reasonably sample the entire colour space boundary. Ideally they would be defined to be close to the solid primary and secondary colours, together with 50% combinations also. The intermediate hues would be defined by all 100%, 75% and 50% combinations of the three chromatic inks, and the neutrals by 25%, 50% and 75% cyan (and the respective amounts of magenta and yellow) respectively. These would be used to generate a test image on the proof and print to ensure that they match to within some tolerance, such as an average ΔE of 2 and a maximum of 5. In fact the information provided by this control target would enable the definition of algorithms, based on Neugebauer or masking equations for example, that define how the press settings should be altered to minimise the error, or the profile modified to eliminate it. Corrections would need to be introduced when the substrate and/or gloss of the proof is different to that of the print.

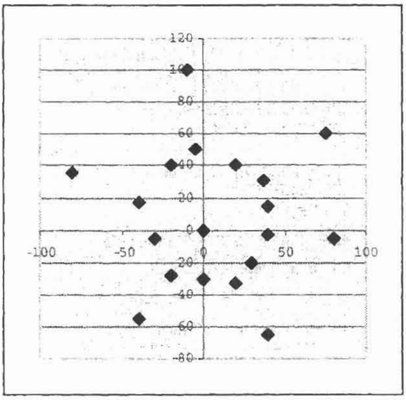

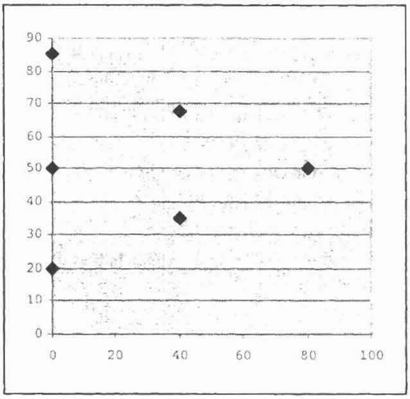

Fig. 3 – a* b* values of control points Fig. 4 – L*, C* values of control points (Note: the above are representative values only – the relative (but not absolute) positions would be retained for each reference gamut. The L*, C* values incorporate 2 hues – those of the primary or secondary colours (the colours with positive C* values and the higher L* values) and the intermediate hue (the colour with a positive C* values and the lower L* value). For completeness figure 4 should really be split into 2 graphs and there should be 6 of each type).

However, ideally what is required for the control element to be effective is that the colours are defined in the exchange encoding space so that they can be processed by the various profiles used in generating proof and print. Unless the encoding space is CMYK data (in which case the control element is clearly easy to define) this creates problems in defining them exactly since the encoding space varies with the scanner – unless it is defined by PCS CIELAB values. Thus there is no unique specification and they have to be defined by a procedure. My recommendation for such a procedure is to select them as the RGB values achieved when the following patches from the target defined in ISO 12641 is scanned: G13-19, K13-19, C16, C4, E4, G4, I4, K4, L4. (Clearly these will not reproduce exactly where figures 3 and 4 imply but they would be reasonably close). The requirement would then be to use these values to produce a test image that would be rendered alongside any images using exactly the same rendering procedures as the image. The CIELAB values achieved on the proof should be measured and specified by the proofer and the test image and CIELAB values should be sent to any printer expected to reproduce it.

Supplementary information

So, the information that needs to be exchanged with an image can be summarised as follows:

- Reference gamut to which the proof was produced (by name and, possibly, encoding).
- Black printer information (at least Total Dot Area, if no 'standard' UCR/GCR can be agreed) where appropriate.
- CIELAB values of control target.
- Control target image.
- Input and output profiles (for existing ICC architecture) or appearance, editing, gamut and input characterisation 'sub-profiles' and associated header information.

Summary

The main topics discussed in this paper are as follows:

- Proposals for a flexible colour management architecture for cross-media reproduction that separates appearance matching, gamut mapping, image editing and characterisation. I believe this will enable greater flexibility, and ensure higher quality and easier optimisation than the current ICC architecture.
- Requirements for the information to be exchanged in addition to the image when a colour match is required.
- Justification for the development of a limited number of reference gamuts for Graphic Arts.
- Black preservation.
- Proposals for procedures for control when using unseparated data for image exchange.

References

Archer, H. B.
1985 "Automatic analysis of reproduction copy for optimum tone reproduction," *TAGA proc. p546*

Braun, G. J. and Fairchild, M. D.
1999 "Image lightness rescaling using sigmoidal contrast enhancement functions," *J. Electronic Imaging, 8*

ANSI Standards
1995 "ANSI CGATS TR 001. Graphic Technology - Colour characterization data for Type 1 Printing" *NPES, Virginia*

Hardy, A. C. and Wurzburg, F.L.
1937 "The theory of three color reproduction," *J. Opt. Soc. Am., 27 p227*
1948 "Color correction in color printing," *J. Opt. Soc. Am., 38 p300*

ICC Specification
1997 "ICC profile format specification - version 3.4" *Current version available from ://www.color.org*

ISO/FDIS 3664
2000 "Viewing conditions for Graphic Technology and Photography" *ISO, Geneva*

ISO 12639
1997 "Graphic Technology - Prepress digital data exchange - Tag image file format for image technology (TIFF/IT)" *ISO, Geneva*

ISO 12641
1996 "Graphic Technology - Prepress digital data exchange - Colour targets for input scanner calibration" *ISO, Geneva*

ISO 12647 (Parts 1 to 6)
1996 "Graphic Technology - Process control for the manufacture of half-tone colour separations, proof and production prints - Part 1, Parameters and measurement methods" *ISO, Geneva*
1996 "Graphic Technology - Process control for the manufacture of half-tone colour separations, proof and production prints - Part 2, Offset processes" *ISO, Geneva*
1999 "FDIS: Graphic Technology - Process control for the manufacture of half-tone colour separations, proof and production prints - Part 3, Coldset offset and letterpress on newsprint" *ISO, Geneva*
200_ "Graphic Technology - Process control for the manufacture of half-tone colour separations, proof and production prints - Part 4, Gravure printing" *ISO, Geneva*
200_ "Graphic Technology - Process control for the manufacture of half-tone colour separations, proof and production prints - Part 5, Screen printing" *ISO, Geneva*
200_ "Graphic Technology - Process control for the manufacture of half-tone colour separations, proof and production prints - Part 6, Flexographic printing" *ISO, Geneva*

ISO 13655
1996 "Graphic Technology - Spectral measurement and colorimetric computation for graphic arts images" *ISO, Geneva*

Johnson, A. J.
1982 "Defining optimum photomechanical colour reproduction," *Pira report PR170*
1985 "Polychromatic colour removal - Evolution or revolution?" *TAGA proc. p1*
1988 "The application of printing specifications to grey component replacement" *TAGA proc. p502*
1989 "Defining optimum photomechanical colour reproduction," *TAGA proc. p350*
1992a "Techniques for reproducing images in different media: advantages and disadvantages of current methods," *TAGA proc., Vol. 2, p739*
1992b "Device independent colour...is it real?," *TAGA proc., Vol. 1*
1995 "A complete colour reproduction model," *TAGA proc., (Vol. 2), p1061*
1997 "Process Control in Colour Reproduction," *TAGA proc. p736*

Johnson, A. J. and Birkenshaw, J. W. B.
1977 "Influence of viewing conditions on colour reproduction objectives," *14th IARIGAI proc. p48*

Johnson, A. J. and Luo, R.
1990 "Colour appearance modelling," *TAGA proc. p144*
1991 "Optimising colour reproduction," *21st IARIGAI proc.*

Jorgensen, G. W.
1976 "Improved black and white half-tones," *GATF Research Project Report No. 105*

McDowell, D. Q.
2000 "Color management – what's needed for printing and publishing," *IPA Bulletin, March/April*

Murray, A. and Morse, R. S.
1941 US Patent 2,253,086

Neugebauer, H. E. J.
1937 "Die theoretischen grundlagen des mehrfarbendruckes," *Z. Wiss. Photogr., 36, 73*

Pobboravsky, I. and Pearson, M.
1972 "Computation of dot areas required to match a colorimetrically specified colour using the modified Neugebauer equations," *TAGA proc. p65*

Schläpfer, K
2000 "Classification of colour gamuts of printing processes," *TAGA proc. (to be published)*

Yule, J. A. C.
1938a "The theory of subtractive colour photography; I, The conditions for perfect colour rendering," *J. Opt. Soc. Am., 28, p419*
1938b "The theory of subtractive colour photography; II, Prediction of errors in colour rendering under given conditions," *J. Opt. Soc. Am., 28, p481*
1967 "Principles of Colour Reproduction," *John Wiley, New York*

Yule, J. A. C. and Colt, R. S.
1951 "Colorimetric investigations in multi-colour printing, " *TAGA proc., p77.*