International Color Consortium®

White Paper #58

Level: Introductory

Date: March 2026

Minimizing the size of output profiles

Introduction

An ICC profile is in effect a container for a colour transform. In many workflows, such as print production, the size of the container file is not critical. However, areas have been identified in which file size is more important, including:

- Exchange of colour images and graphics on the web
- Embedding profiles in PDF/A [1] files
- Variable data printing workflows

Some ICC profile classes have a compact representation of the transform between device space and the Profile Connection Space (PCS). The older 'Matrix/TRC' profiles can encode just a 3x3 matrix and parameters for a set of one-dimensional curves, while v4 LUT-based profiles can also encode just matrix and 1-D A curves, leaving the other curves and colour look-up table empty (unless the number of channels in the device data is not equal to 3). The new CICP tag enables the profile to simply carry code points pointing to transform parameters in ISO/IEC 23091-2 [2], while in iccMAX it is also possible to encode a rich set of encoding parameters in a 'cenc' profile [3] by pointing to external colour encodings. All these methods can yield a profile of around 1Kb in size.

Output class profiles, on the other hand, do not have similarly compact representations. The transform is stored in the profile in a tag which includes a multi-dimensional look-up table encoded as 16-bit integers with num_grid_points ^ num_channels entries, as well as one-dimensional curves and matrices; and the format specification mandates that a tag for each of the three rendering intents in both source-to-PCS and PCS-to-destination direction is present. As a result the average size of CMYK profile in the ICC Profile Registry is 3.45MB (with a range of 0.7 to 8.4MB).

Embedding such profiles in media files can lead to significantly larger file sizes, which can pose challenges for data transmission over networks and impact web performance.

The profile format is intended to make it possible to encode a high quality colour space transform with minimal quality losses which might lead to colorimetric errors or artefacts such as contouring. When attempting to reduce file size of output profiles, the key question is what adjustments can be made to profile creation that lead to minimal losses in quality, and what is the trade-off between profile size and transform quality.

The criteria for assessing profile quality are primarily colorimetric accuracy (in the Colorimetric rendering intents) and pleasingness (in the Perceptual intent). To these can be added a lack of contouring, caused by steep changes in the transform. This is important in both Colorimetric and Perceptual intents.

Size of elements

The file size of an output profile is mainly comprised of the elements shown in Table 1.

TABLE 1: TYPICAL SIZE OF PROFILE ELEMENTS

Element	**Fixed or variable**	**Typical size (kB)**
Header	Fixed	0.128
Tag table	Fixed	0.012 per tag
Essential metadata (copyright and profile description, white point and chad tags, typically around 300 bytes)	Variable	0.3
Optional metadata such as characterization data from which the profile was made	Variable	126
Gamut tag	Variable	33
AToB LUTs	Variable	3 x 400
BToA LUTs	Variable	3 x 300
Private tags	Variable	N/A
Total (without private tags)	Variable	2260

A strategy to minimize the profile size would aim to include only the elements mandated by the specification (thus avoiding the use of optional and private tags), and reducing the size of the required tags, especially the largest ones, to the smallest possible consistent with the quality requirements.

Reducing the size of LUTs

The ICC specification does not support data compression, so the LUT elements can be considered for size reduction as follows.

Precision

There is a choice of 8 or 16-bit integer data types for the curves and the multi-dimensional color look-up table (CLUT). Using 16-bit data will of course double the size of the elements.

Unique intents

The ICC specification requires that the tag table of an output profile includes the byte offset and size of all three rendering intents. However, it allows multiple intents to have the same offset, so that a single AToB tag can be used for two or even all three required intents; and similarly a single BToA tag can stand for two or more intents. The Saturation intent is in practice not always required or encoded as a unique transform, in which case the tag table can point to the byte offset of the Perceptual intent. If size reduction is especially important, the Perceptual and Saturation tags could use the data for the Colorimetric intent, although if this is done the Perceptual intent in a v4 profile will not use the Perceptual Rendering Medium and will be incorrect. Alternatively, to ensure conformance, where it is envisaged that a reduced-size profile is intended for use with a single rendering intent, the alternate intents can be populated with minimum-size CLUTs as described below.

Curves

ICC curve types can be single-exponent (in which a single gamma value is encoded for each curve), parametric (encoding a set of parameters according to the set of supported parametric curve types) or tabulated (encoding a sample of points along the curve which can be interpolated by the CMM). Tabulated curves are the largest, requiring one or two bytes for each point in the curve, depending on whether 8- or 16-bit precision is used.

In creating a profile the choice is often made to linearise the output device prior to the PCS using a tabulated A curve, and common sizes include 256 and 1012 points. Many fewer points can be used without loss of accuracy; for output devices the sampling precision should be equal to or less than the data precision of the output channels (e.g. 256 points for 8-bit data).

If a CMM is using linear interpolation, using a very small number of points to define the curve could possibly lead to perceptible errors, such as contouring, if the curve ihas flat segments or large slope changes.

Possible strategies for reducing the size of the curve elements are to

- Reduce the number of points used to define the curve. As indicated above, this could lead to a degree of contouring even if the curve is a close fit to the data, due to CMM interpolation.

- Replace the curve with a null element or a straight line curve, defined just by 0 and 65535 start and end points, and handling the linearization within the CLUT.

However, a 1-D curve is much smaller than a multi-dimensional CLUT, and the more optimal tradeoff may be to use the A curves to linearize the device data with respect to the PCS in order to minimize the number of grid points needed in the CLUT. Experience shows that 21 uniformly-spaced curve points is often sufficient, while 42 points captures almost all of the points in the single-ink ramps in an ISO 12642 test chart.

CLUT

The CLUT is normally the largest single element in an output profile, and its size is determined by the number of channels in the device space, the number of grid points in each dimension, and the data precision. The size of the CLUT element in bytes is num_gridpoints^num_channels*precision, where num_channels is the number of channels in the PCS (for an AToB tag) or the data colour space (for a BToA tag), and the precision is the number of bytes per entry, either one or two. Hence the number of grid points in each dimension is critical to CLUT size.

The specification permits up to 255 grid points per dimension, although in practice profiles with more than 33 are rarely seen. The minimum is two, which would represent a linear interpolation of the data. Small profiles can be made with, for example, 5 or 9 grid points per dimension, with a potential reduction in accuracy since the CMM is interpolating across larger sub-regions of colour space which may be non-linear. In an ICC v4 or iccMAX profile the number of grid points does not have to be the same in each dimension, although in practice they commonly are.

A high LUT density minimizes errors when the output space is non-linear with respect to the input, but if linearity is already good, or if linearization can be achieved in the curves, CLUT sampling can be reduced without loss of accuracy.

Subtractive colorant channels cyan, magenta, yellow and black (and additive channels red, green and blue) are in practice often smooth and close to linear with respect to CIELAB L*, making it possible to reduce curve density in A curves.

Modifying the ratio of grid points in each dimension could enable some optimization strategies, but could also make the profile creation more complex. CIELAB a* and b* channels can be reduced more in density than the L* channel before errors become perceptible.

In certain cases a profile may be intended primarily for use in one direction only. In such cases the number of grid points in the BToA or AToB tag which converts in the opposite direction can be further reduced. In extreme cases just two grid points in each dimension can be encoded, although this is not recommended in cases where a profile could conceivably be expected to give acceptable results in the inverse direction, including PDF and web.

Gamut tag

The gamut tag is a required element in the profile, but is not essential to CMM processing. The tag includes a CLUT and when file sizes are critical this can be encoded as just two grid points in each dimension. A gamut representation thus encoded will of course not represent the gamut of the output data colour space medium.

Effect on profile quality

If every option to reduce the size of elements in a profile is taken, a CMYK output profile of under 50Kb is possible. That being the case, it is important to understand the consequences of such a small profile on the criteria for profile quality.

Students at NTNU conducted a study to generate CMYK profiles with a small footprint and evaluate their quality. Since the Adobe FOGRA39 profile is a common choice in PDF/A workflows, this was used as to benchmark results against. The FOGRA39 data set from the Characterization Data Registry was selected, and Argyll

CMS was chosen for profile creation as it is open source software that handles colour separation, black generation, gamut mapping and so on, which were not the focus of the study.

A range of options for file size reduction were progressively compared to identify the effect on the colorimetric performance of the profile, and a psychophysical study was conducted to explore the visual impact on reproduction quality. Editing the profile was performed on an XML representation after converting with IccXml.

The steps in generating the small-footprint profile were:

1. Create profile in Argyll CMS with 'medium quality' option – 9 grid points per channel in both AToB and BToA tags
2. Delete all non-required tags
3. Delete Perceptual and Saturation intents and edit tag table so that the byte offsets for three intents point to the Colorimetric intent
4. Replace each curve with 2 points for linear curves or 8 points for non-linear curves
5. Convert 16-bit LUT and curve data to 8-bit

Table 2 shows the independent effect on the file size of each step.

TABLE 2: EFFECT OF EACH REDUCTION STEP ON RELATIVE FILE SIZE

	File size (KB)	Reduction (KB)
Base profile	192	0
Non-required tags removed	116	76
Linear curves	186	6
8-point curves	162	30
8-bit data	123	69

Colorimetric performance

The profiles were compared with the Adobe FOGRA39 profile in both AToB and BToA directions, using a test chart. Table 2 shows the independent effect on the colorimetric performance.

TABLE 3: EFFECT OF FILE SIZE REDUCTION ON RELATIVE ACCURACY OF PROFILE

	AToB		BToA	
	ΔE^*_{ab} avg	ΔE^*_{ab} max	ΔE^*_{ab} avg	ΔE^*_{ab} max
Base profile	0.72	3.45	2.22	5.0
Non-required tags removed	0.72	3.45	2.22	5.0

Linear curves	0.72	3.45	2.22	5.0
8-point curves	0.78	3.46	2.25	5.0
8-bit data	1.28	4.12	2.44	5.0

A heat map of differences between the Adobe FOGRA39 profile and the 8-bit small-footprint profile is shown in Figure 1

FIGURE 1: **HEAT MAP ILLUSTRATING COLORIMETRIC DIFFERENCES BETWEEN PROFILES**

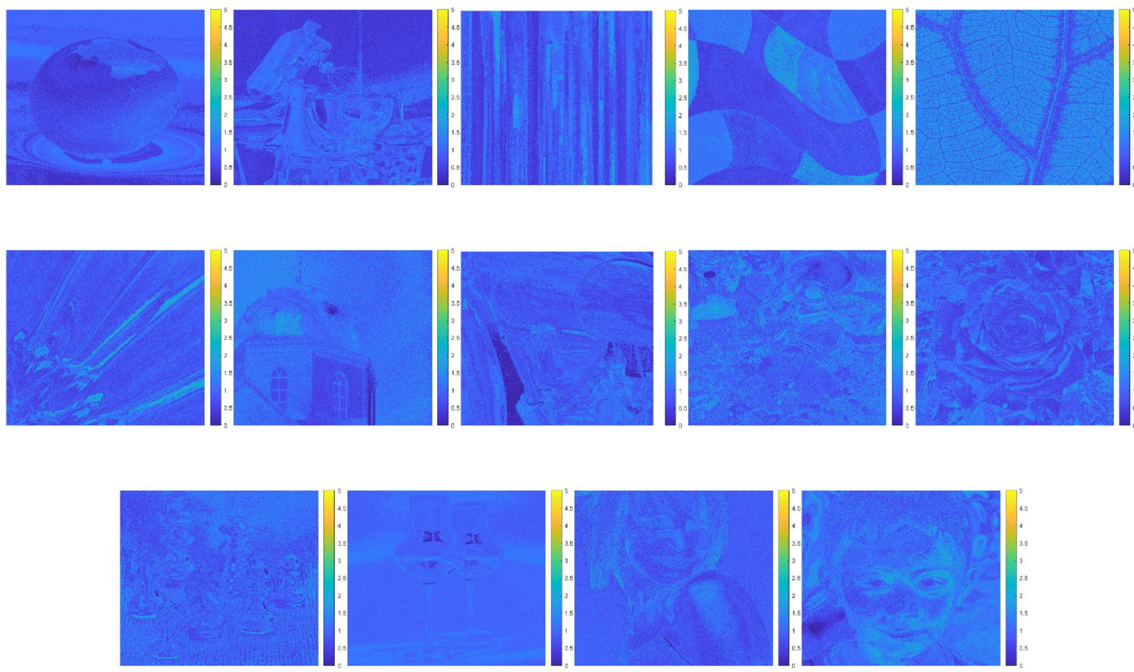

Visual performance

A psychophysical test was performed in which small-footprint profiles with 8- and 16-bit LUTs were compared with the Adobe FOGRA39 profile. Seven test images were transformed to CMYK using the different profiles, and the match between Adobe FOGRA 39 and the small-footprint profiles was judged by 25 observers in a two-alternative forced-choice experiment, in both AToB and BToA directions. The small-footprint conversions were judged to be indistinguishable to those of the Adobe FOGRA 39 profile in the AToB direction for both 8- and 16-bit LUTs, while in the BToA direction differences in some images were visible for both LUT sizes.

Smoothness test

Since it might be expected that a LUT with fewer grid points would lead to a lack of smoothness in the output when the input is a smooth and continuous transition, a smoothness test was performed. The smoothness metric described in [4] was applied; test images from [4] were converted from CIELAB to CMYK using the small-footprint profile, and the sum of second derivatives along each of the test images calculated. The results are shown in Table 4, and compared with results of the same test applied to the original profile. It can be seen that the reduction in smoothness is negligible.

FIGURE 2: SMOOTHNESS TEST IMAGES, WITH CONTINUOUS GRADIENTS IN CIELAB SPACE (CONVERTED TO DOCUMENT SPACE RGB FOR ILLUSTRATION)

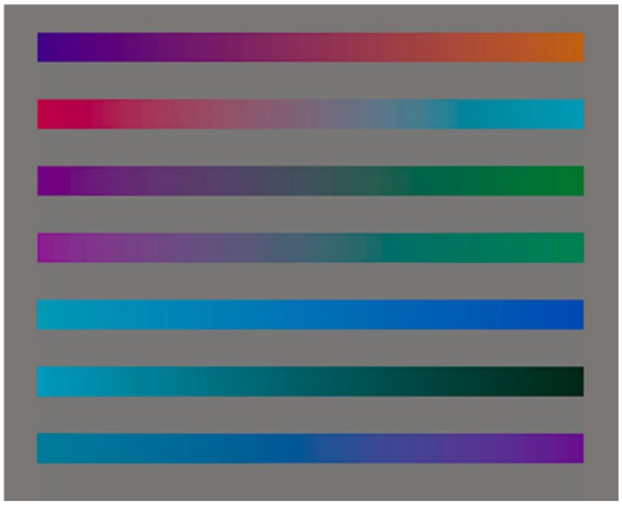

FIGURE 3: SMALL AREA OF TEST IMAGE AFTER CONVERSION TO CMYK: A) ADOBE FOGRA39 PROFILE; B) 8-BIT SMALL-FOOTPRINT PROFILE (CONVERTED TO DOCUMENT SPACE RGB FOR ILLUSTRATION)

TABLE 4: SMOOTHNESS METRIC RESULTS FOR THE 7 TEST IMAGES, AFTER CONVERTING FROM CIELAB TO CMYK

Test image no		1	2	3	4	5	6	7
	L*	0	0	0	0	0	0	0
Original image	a*	55	75	57	0	0	0	0
	b*	0	0	0	0	0	0	0
	L*	0	0	0	0	0	0	0
Adobe FOGRA39 profile	a*	49	63	42	0	0	0	0
	b*	0	0	0	0	0	0	0
	L*	0	0	0	0	0	0	0
2024 8-bit small-footprint	a*	52	67	44	0	0	0	0
	b*	0	0	0	0	0	0	0

Compression

At this time ICC makes no recommendation on profile compression. However, ISO/IEC 18181-1 [5] (defines a lossless compression scheme for byte-stream encoded ICC profiles in JPEG XL images using a Brotli custom dictionary. This method is under evaluation by ICC and may form a future recommendation.

Conclusion

CMYK output profiles tend to have large file sizes in order to support the encoding of accurate conversions between the PCS and CMYK, with the largest elements being in the multi-dimensional LUTs. Where smaller file sizes are needed it is possible to remove or reduce the size of elements with minimal effect on colorimetric accuracy or smoothness. The largest file size reductions can be achieved by using a single AToBx and BToAx tag for all three rendering intents, reducing the number of grid points in each dimension, removing any non-required tags, using 8-bit instead of 16-bit data, and reducing the number of samples in tabulated curves. It should however be noted that a v4 ICC profile with only a single unique rendering intent would violate the Perceptual Reference Medium requirements for the Perceptual intent.

References

1. ISO 19005 (all parts), International Organization for Standardization. Document management – Electronic document file format for long-term preservation
2. ISO/IEC 23091-2:2025 Information technology — Coding-independent code points Part 2: Video, ISO, Geneva
3. ISO 20677:2019 Image technology colour management — Extensions to architecture, profile format and data structure
4. P, Green, A smoothness metric for colour transforms. Proc SPIE. 6807 (2008)
5. ISO/IEC 18181-1:2024 Information technology — JPEG XL image coding system Part 1: Core coding system, ISO, Genev