

Deferred Demosaicing of an ARRIRAW Image File to a Wide- Gamut Logarithmic Encoding



Page 1 of 17 pages

The attached document is a Registered Disclosure Document prepared by the proponent identified below. It has been examined by the appropriate SMPTE Technology Committee and is believed to contain adequate information to satisfy the objectives defined in the Scope, and to be technically consistent.

This document is NOT a Standard, Recommended Practice or Engineering Guideline, and does NOT imply a finding or representation of the Society.

Errors in this document should be reported to the proponent identified below, with a copy to eng@smpte.org.

This document is intended to support the development of applications that read and process ARRIRAW image files. It is not intended to support the development of hardware or software applications that create ARRIRAW image files, and creation of such files is reserved to individuals and organizations that have entered into agreements with the proponent identified below for such file creation. All inquiries in respect of this document (other than reporting of errors in the document, which should be handled as described in the prior paragraph), including inquiries as to intellectual property requirements that may be attached to use of the disclosed technology, should be addressed to the proponent identified below.

Proponent contact information:

Joseph Goldstone
Arnold & Richter Cine Technik GmbH & Co. Betriebs (KG)
Turkenstraße 89
D-80799 München
Germany
Email: jgoldstone@arri.com

Table of Contents	Page
Introduction	3
1 Scope.....	4
2 Normative References	4
3 Acronyms.....	4
4 Trademarks.....	5
5 Definitions of the Algorithms.....	5
5.1 Unpacking and Linearization of Near-Logarithmic Unsigned Integer Data from ARRIRAW File...	5
5.2 Application of Per-Channel White Balance Factors	8
5.3 Demosaicing of the Photosite Data	9
5.4 Conversion to Wide-Gamut Encoding	9
5.5 Exposure Compensation	13
5.6 Logarithmic Encoding	13
Annex A Bibliography (Informative)	14
Annex B Color Rendering Example (Informative)	15
B.1 Tonemapping.....	15
B.2 Conversion for Displays with Specific Primaries, White Points and EOTFs	16

Introduction

This document describes the deferred demosaicing and logarithmic encoding of raw image data stored in ARRIRAW files produced by ARRI cameras, as described in SMPTE RDD 30, ARRIRAW Image File Structure and Interpretation Supporting Deferred Demosaicing to a Logarithmic Encoding, producing a sequence of image data as RGB triplets representing a logarithmic encoding of exposure values, in a color space whose primaries are set well enough apart to encompass the vast majority of real-world scene colors, and with a clearly-defined neutral chromaticity. These resulting image data represent successive raster lines of an image whose vertical and horizontal extents match those recorded in the SMPTE RDD 30 image file header.

There are many possible uses for this logarithmically encoded raster data, including:

- Color rendering for display. An example is given as an informative Annex to this document in which the logarithmic data are tone-mapped, re-encoded for the primaries and white point for a common class of display (using a transform that incorporates chromatic adaptation), and remapped to compensate for the presumed EOTF of that same class of display.
- Placement (with accompanying metadata) into a container for later processing, cf. SMPTE ST 268 (SMPTE Standard for File Format for Digital Moving-Picture Exchange (DPX), Version 2.0, and Codex Technical Bulletin — Codex DPX Multichunk User Data, Revision 15.10.
- Processing by an ACES device transform to produce an ACES image, cf. SMPTE ST 2065-1 (Academy Color Encoding Specification) and P-2013-001 (Recommended Procedures for the Creation and Use of Digital Camera System Input Device Transforms (IDTs)). Such an image might be subject to immediate manipulation, or might be stored into a container for later processing, as described in SMPTE ST 2065-4 (ACES Image Container File Layout).

The demosaicing process is preceded by application of per-channel white balance factors to the radiometrically linear photosite data, and followed by exposure compensation, the application of a 3x4 matrix that transforms reconstructed sensor RGB values to a color encoding with a defined white point and a set of RGB primaries, and finally a logarithmic encoding function that is dependent on exposure index.

The channel-dependent white balance factors are taken from the image content information subheader of the raw image file. The 3x4 matrix that converts reconstructed sensor RGB values to the RGB color space defined for the logarithmic encoding may be stored in that same image content information subheader, or if not, the matrix may be reconstructed from information in this document appropriate for the exposure index given in that same subheader. The formula for exposure compensation is not carried as metadata, nor is any equation explicitly carried as metadata that would effect the logarithmic encoding; both of these, however, are given in this document.

It is the intent of this RDD to document the structure and interpretation of ARRIRAW files generated by ARRI cameras, so that users of this document may develop applications correctly identifying and interpreting such files. This document is specifically not intended to support development of hardware or of software applications creating ARRIRAW files. Permission to create such files, along with additional documentation to support that creation, is reserved by ARRI to members of the ARRI Partner Program, the contact information for which is provided below.

Assistance in correctly processing ARRIRAW files, including certification that the results of processing meet ARRI quality standards, is available to members of the ARRI Partner Program. Information on the ARRI Partner Program is available at the address given below:

Digital Workflow Solutions group – ARRI Partner Program
 Arnold & Richter Cine Technik GmbH & Co. Betriebs (KG)
 Turkenstraße 89
 D-80799 München
 Germany
 Email: dws@arri.de

1 Scope

This document describes the deferred demosaicing and logarithmic encoding of ARRIRAW image data stored as described in SMPTE RDD 30 to produce a sequence of encoded image data represented by ARRI Wide Gamut RGB values

2 Normative References

Note: All references in this document to other SMPTE documents use the current numbering style (e.g. SMPTE ST 268:2003) although, during a transitional phase, the document as published (printed or PDF) may bear an older designation (such as SMPTE 268M-2003). Documents with the same root number (e.g. 268) and publication year (e.g. 2003) are functionally identical.

The following standards contain provisions that, through reference in this text, constitute provisions of this registered disclosure document. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this registered disclosure document are encouraged to investigate the possibility of applying the most recent edition of the standards indicated below.

CIE 15:2004, 3rd Edition, Colorimetry.

ISO 11664-1:2007 (CIE S 014-1/E:2006), Colorimetry — Part 1: CIE Standard Colorimetric Observers

ISO 12232:2006, Photography — Digital Still Cameras — Determination of Exposure Index, ISO Speed Ratings, Standard Output Sensitivity, and Recommended Exposure Index

Recommendation ITU-R BT.709 (04/2002), Parameter Values for the HDTV Standards for Production and International Programme Exchange

Recommendation ITU-R BT.1886 (03/2011), Reference Electro-Optical Transfer Function for Flat Panel Displays Used in HDTV Studio Production

SMPTE RDD 30:2014, ARRIRAW Image File Structure and Interpretation Supporting Deferred Demosaicing to a Logarithmic Encoding

SMPTE ST 268:2003, File Format for Digital Moving-Picture Exchange (DPX), Version 2.0

Amendment 1:2012 to SMPTE ST 268:2003

SMPTE ST 2065-1:2012, Academy Color Encoding Specification (ACES)

SMPTE ST 2065-4:2013, ACES Image Container File Layout

SMPTE RP 431-2:2011, D-Cinema Quality — Reference Projector and Environment

3 Acronyms

Acronyms used in this document are listed below.

ACES: Academy Color Encoding Specification

CCT: Correlated Color Temperature

DPX: Digital Picture Exchange

EOTF: Electro-Optical Transfer Function

IDT: Input Device Transform

IR: Infrared

LDS: Lens Data System

N/A: Not Applicable

ND: Neutral Density

UUID: Universally Unique Identifier

UV: Ultraviolet

VFX: Visual Effects

4 Trademarks

ALEXA is a registered trademark of Arnold & Richter Cine Technik GmbH & Co.

5 Definitions of the Algorithms

The process of converting image data and image metadata in an ARRIRAW file into a logarithmically encoded image is divided into six steps:

- Unpacking and linearization of near-logarithmic 12-bit unsigned integer encoded photosite values in the ARRIRAW file to produce linear 16-bit unsigned integer photosite values;
- Application of per-channel white balance factors to the linear 16-bit unsigned integer photosite data, with the factors being dependent on correlated color temperature and a magenta/green bias factor;
- Demosaicing of the white-balanced photosite data;
- Conversion of the reconstructed sensor RGB by a 3x3 matrix, with the matrix coefficients being dependent on correlated color temperature and the presence or absence of an internal camera filter, to ARRI Wide Gamut RGB, a wide-gamut RGB encoding with defined primaries and white point;
- Exposure compensation, with the amount of compensation being dependent on exposure index; and finally
- Logarithmic encoding, with scaling and offset factors of the encoding algorithm determined by exposure index.

Each of these steps is described below.

5.1 Unpacking and Linearization of Near-Logarithmic Unsigned Integer Data from ARRIRAW File

Every sequential block of 12 8-bit bytes corresponds to 3 sequential 32-bit words containing 8 near-logarithmic 12-bit unsigned integer photosite values. These 8 near-logarithmic 12-bit unsigned integer photosite values are unpacked to form 8 linear 16-bit unsigned integer photosite values using an algorithm such as (or equivalent to) that given in Section 5.1.2.1.2.


5.1.1 Reading of Image File Data as Packed Data

The packed data are read as 8-bit bytes, every successive 4 of which are to be assembled into a 32-bit word.

Each byte contains photosite data in two half-bytes each comprised of 4 bits. Such a 4-bit half-byte represents either the most significant 4 bits of a 12-bit near-logarithmic value, the next-most-significant 4 bits of a 12-bit near-logarithmic value, or the least significant 4 bits of a 12-bit near-logarithmic value. In the figure, these three interpretations of a 4-bit half-byte are represented by a large trapezoid, a small trapezoid, and a triangle, respectively.




In the assembly of a 32-bit word from 4 8-bit bytes, the least significant 8 bits are read first, and the most significant 8 bits are read last. The resulting 3 32-bit words represent 8 successive packed photosite values, suitable for input to the algorithm given in pseudocode in Section 5.1.2.




5.1.2 Unpacking of 32-bit Words into 12-bit Encoded Values

The 12-bit quantities that have been packed into 32-bit words are unpacked prior to processing. The algorithm whose pseudocode is given in Section 5.1.2.1.2 performs this unpacking. It assumes that successive 32-bit words are stored adjacently in memory, and that the unpacking algorithm is executed on a machine with a little-endian architecture, taking as input blocks of 3 32-bit words and producing as output blocks of 4 32-bit words. In the figure below, the top element represents the unpacked data addressed as 32-bit words and the bottom element represents the unpacked data addressed as 16-bit half-words.

Every resulting 32-bit word of unpacked data contains two 12-bit near-logarithmic photosite values. The lower-order 12 bits of the least significant 16 bits of the 32-bit word contain the 12-bit near-logarithmic value corresponding to the photosite encountered earlier in sensor traversal; the lower-order 12 bits of the most significant 16 bits of the 32-bit word contain the 12-bit near-logarithmic value encountered later in sensor traversal.



When addressed as 16-bit half-words, those same unpacked data are in sensor traversal order and are interpreted as unsigned integers representing 12-bit near-logarithmic encoded photosite values.



5.1.2.1.1 Variable definitions

NumPixel	number of pixels in the unpacked image
InSize	memory used to store the 12-bit packed image
OutSize	memory size required to store the unpacked image

lpIn	pointer to the input packed image memory
lpOut	pointer to the output unpacked image memory
MSB_Pixel	pixel in the 16 most significant bits of the output 32-bit memory word
LSB_Pixel	pixel in the 16 least significant bits of the output 32-bit memory word
InWord	input 32-bit packed word

5.1.2.1.2 Pseudocode implementation

```

void ARI_UnPack12BitBayerPattern(uint32 *lpIn, uint32 *lpOut, ...)
{
    OutSize = 16*InSize/12;
    NumPixel = OutSize * 2;
    While (NumPixel > 0)
    {
        InWord = lpIn[0];
        MSB_Pixel = ((InWord >> 20) & 0x00000FFF);
        LSB_Pixel = ((InWord >> 8) & 0x00000FFF);
        lpOut[0] = ((MSB_Pixel << 16) & 0xFFFF0000) | LSB_Pixel;
        MSB_Pixel = ((InWord << 4) & 0x00000FF0);
        InWord = lpIn[1];
        MSB_Pixel |= ((InWord >> 28) & 0x0000000F);
        LSB_Pixel = ((InWord >> 16) & 0x00000FFF);
        lpOut[1] = ((MSB_Pixel << 16) & 0xFFFF0000) | LSB_Pixel;
        MSB_Pixel = ((InWord >> 4) & 0x00000FFF);
        LSB_Pixel = ((InWord << 8) & 0x00000F00);
        InWord = lpIn[2];
        LSB_Pixel |= ((InWord >> 24) & 0x000000FF);
        lpOut[2] = ((MSB_Pixel << 16) & 0xFFFF0000) | LSB_Pixel;
        MSB_Pixel = ((InWord >> 12) & 0x00000FFF);
        LSB_Pixel = (InWord & 0x00000FFF);
        lpOut[3] = ((MSB_Pixel << 16) & 0xFFFF0000) | LSB_Pixel;
        lpIn += 3;
    }
}

```

```

        lpOut += 4;
        NumPixel -= 8;
    }
}

```

5.1.3 Linearization of 12-bit Near-Logarithmic Unpacked Sensor Data

Photosite values are radiometrically linear representations of the energy they receive, represented as 16-bit unsigned integers, incorporating an offset such that a photosite receiving no energy would have a 16-bit unsigned integer value of 256.

When stored as ARRIRAW file image data, the 16-bit photosite values are encoded into 12-bit values using a method that approximates a logarithmic encoding, but emphasizes the preservation of deep shadow detail.

The following equation reconstructs a 16-bit linear photosite value v_p from an encoded 12-bit value v_i .

$$v_p = \begin{cases} v_i, & v_i < 1024 \\ (((1024 + 2 \cdot o + 1) \ll (q - 2)) - 1), & v_i \geq 1024 \end{cases}$$

where

q is the integer part of the quotient of v_i divided by 512

o is v_i modulo 512

5.2 Application of Per-Channel White Balance Factors

After first subtracting the camera's average black level (which is fixed at value 256 for all files whose processing is described by this document), the linear 16-bit unsigned integer photosite data are multiplied by the White Balance Factors in the Image Content Information subheader of the ARRIRAW file that is appropriate to their type, after which the average black level is added back to the calculated products. The mapping of photosite type to corresponding White Balance Factor is as follows:

Photosite Type	Corresponding Factor
G ₀	White Balance Green Factor
R	White Balance Red Factor
B	White Balance Blue Factor
G ₁	White Balance Green Factor

The White Balance Green Factor is unity:

$$g0_{balanced} = g0_{photosite}$$

$$g1_{balanced} = g1_{photosite}$$

When the 16-bit unsigned integer values representing photosite data are converted to equivalent floating-point values, and the White Balance Factors R and B are applied using floating-point arithmetic to these floating-point red and blue photosite data values, the calculations can each be performed with three operations:

$$r_{balanced} = (r_{photosite} - 256) * R + 256$$

$$b_{balanced} = (b_{photosite} - 256) * B + 256$$

The black-subtracted intermediate results in each case are converted to floating-point prior to multiplication by the floating-point white balance factors, with the resulting product converted back to 16-bit integer form prior to the addition of the previously-subtracted black level.

When the White Balance Factors R and B are applied to red and green photosite data values using unsigned integer arithmetic, careful ordering of operations should be used to ensure that the lower extent of black noise (i.e. photosite values below 256) are not clipped.

Three example sets of White Balance Factor Values are given below, with the Green/Magenta Tint presumed to be zero.

CCT	Tint	White Balance Red Factor	White Balance Green Factor	White Balance Blue Factor
3200	0	1.128195	1	2.068762
5600	0	1.644962	1	1.366723
6500	0	1.745829	1	1.252820

At time of capture these factors are calculated from the chosen White Balance CCT and Green/Magenta Tint values.

5.3 Demosaicing of the Photosite Data

A demosaicing algorithm is applied to the photosite data to produce spatially correspondent RGB data. As the photosites are arranged in a classic Bayer pattern, there are many algorithms in the literature. For an introduction, see “Color Processing in Digital Cameras”, by Adams, Parulski and Spaulding, in IEEE Micro, vol. 18 (6), pp. 20-30. For a more varied treatment see the compilation of papers in “Single Sensor Imaging”, Rastislav Lukac, Ed., CRC Press 2009.

5.4 Conversion to Wide-Gamut Encoding

The reconstructed sensor RGB data are transformed to the ARRI Wide Gamut color space, a wide-gamut color space the primaries and white point of which are defined by the CIE x,y chromaticity coordinates given below.

	x	y
Red	0.6840	0.3130
Green	0.2210	0.8480
Blue	0.0861	-0.1020
White	0.3127	0.3290

A 3x3 matrix M transforms sensor RGB to ARRI Wide Gamut RGB. The application of this matrix is performed on the result of a subtraction of the camera’s average black level from the sensor RGB data, and the camera’s average black level is added to the result of the matrix application. When used in floating-point arithmetic calculations, this matrix is applied to the left of a column vector of sensor RGB whose values have

been reduced by the camera's average black level. The results of the matrix application are then augmented by the camera's average black level.

$$RGB_{wg} = RGB_{black} + M \cdot (RGB_{sensor} - RGB_{black})$$

When such floating-point calculations are performed the value of each of the three components of RGB_{black} is fixed at value 256 for all files whose processing is described by this document.

The black-subtracted intermediate results are converted to floating-point prior to multiplication by the floating-point matrix, with the resulting product converted back to 16-bit integer form prior to the addition of the previously-subtracted black level.

When matrix operations are applied to sensor RGB using unsigned integer arithmetic, careful ordering of operations should be used to ensure that the lower extent of black noise is not clipped.

The contents of a 3x3 matrix M appropriate for the correlated color temperature recorded in the White Balance CCT field of the Image Data Information subheader are carried in metadata as the first three columns of the Color Matrix field of that same subheader.

If those Color Matrix metadata fields have not been filled in, the following two tables give matrix coefficients for M appropriate to the correlated color temperature recorded in the White Balance CCT field of that same Image Data Information subheader. The choice of which table is used depends on the content of the ND Filter field of the Lens Data Information Subheader of the ARRIRAW image metadata.

White Balance CCT values within the range described below, but falling between the listed CCTs, should use matrices that are linearly interpolated between the immediately surrounding rows. A single interpolation parameter a controls this interpolation, with an interpolated matrix M being computed as

$$M = (1 - a) M_1 + a M_2$$

where

the interpolation factor a for CCT c being computed as

$$a = \frac{\frac{1}{c} - \frac{1}{CCT_1}}{\frac{1}{CCT_2} - \frac{1}{CCT_1}}$$

M_1 and CCT_1 are the matrix and CCT from one adjacent row

M_2 and CCT_2 are the matrix and CCT from the other adjacent row

The CCT values CCT_1 and CCT_2 should be converted to floating-point numbers prior to their reciprocal values being computed, and all subsequent mathematical operations computing the interpolated matrix M should be performed using floating-point arithmetic.

These matrices are derived from the camera system's measured spectral response, from a presumed illuminant, and from data sets of real-world spectral reflectances. The sensor response of real-world cameras being invariably different from those of the CIE 1931 Color Matching Functions, such a derivation process for a 3x3 matrix always involves error minimization, and possibly may involve shifting residual error towards colors less frequently encountered in real-world scenes. Such error can be moved to different parts of the space of possible sensor RGB values but it cannot be completely eliminated. It is therefore possible for sensor RGB at the periphery of sensor RGB space to be transformed into invalid colors (e.g. with negative R, G and/or B components, given the negative off-axis matrix elements) in the ARRI Wide Gamut color space. Gamut mapping or other methods for the handling of such invalid colors are outside the scope of this document.

5.4.1 Conversion Matrix Contents when no ALEXA Studio ND Type 1 Filter was used

When the Filter Type component of the ND Filter field of the Lens Data Information subheader indicates no ND filter was in place at time of acquisition, the following table contains the matrix coefficients that are used for several common CCTs.

CCT	$M_{1,1}$	$M_{1,2}$	$M_{1,3}$	$M_{2,1}$	$M_{2,2}$	$M_{2,3}$	$M_{3,1}$	$M_{3,2}$	$M_{3,3}$
2000	1.210510	-0.262282	0.051773	-0.121371	1.051117	0.070254	0.001944	-0.300355	1.298410
2100	1.202154	-0.252435	0.050281	-0.110522	1.039042	0.071480	0.007013	-0.298648	1.291635
2200	1.195132	-0.243392	0.048261	-0.101274	1.030202	0.071073	0.010983	-0.295148	1.284164
2400	1.184217	-0.227449	0.043233	-0.086385	1.019161	0.067224	0.016703	-0.285655	1.268952
2600	1.176497	-0.213934	0.037437	-0.074943	1.013870	0.061073	0.020564	-0.275263	1.254699
2900	1.169237	-0.197282	0.028045	-0.062025	1.012171	0.049854	0.024440	-0.260551	1.236111
3200	1.165689	-0.184001	0.018311	-0.052436	1.014726	0.037709	0.027069	-0.247829	1.220760
3500	1.164751	-0.173288	0.008537	-0.045025	1.019609	0.025416	0.029018	-0.237114	1.208095
3900	1.166292	-0.162017	-0.004275	-0.037402	1.028012	0.009390	0.031011	-0.225453	1.194443
4300	1.162697	-0.142148	-0.020548	-0.032953	1.031553	0.001399	0.032608	-0.224601	1.191993
4700	1.165366	-0.132799	-0.032567	-0.027832	1.040165	-0.012334	0.034112	-0.216134	1.182022
5100	1.169736	-0.125665	-0.044071	-0.023631	1.049208	-0.025577	0.035391	-0.209090	1.173699
5600	1.176639	-0.119021	-0.057618	-0.019367	1.060570	-0.041202	0.036749	-0.201922	1.165172
6500	1.190760	-0.111827	-0.078933	-0.013774	1.079540	-0.065767	0.038524	-0.192614	1.154090
7500	1.206597	-0.106314	-0.100283	-0.009227	1.099353	-0.090126	0.040350	-0.185178	1.144828
9000	1.227895	-0.102511	-0.125383	-0.004820	1.123352	-0.118532	0.042160	-0.178176	1.136016
11000	1.250764	-0.100372	-0.150393	-0.001199	1.147714	-0.146515	0.043813	-0.172586	1.128773

5.4.2 Conversion Matrix Contents when an ALEXA Studio ND Type 1 Filter was used

When the Filter Type component of the ND Filter field of the Lens Data Information subheader indicates that an ALEXA Studio ND Type 1 filter was in place at time of acquisition, the following table contains the matrix coefficients that are used for several common CCTs.

CCT	M _{1,1}	M _{1,2}	M _{1,3}	M _{2,1}	M _{2,2}	M _{2,3}	M _{3,1}	M _{3,2}	M _{3,3}
2000	1.158737	-0.169501	0.010764	-0.141800	1.033172	0.108628	-0.007010	-0.482172	1.489182
2100	1.149670	-0.162385	0.012715	-0.127104	1.016543	0.110561	0.006045	-0.481253	1.475208
2200	1.141972	-0.155434	0.013462	-0.114787	1.004343	0.110444	0.016390	-0.477733	1.461343
2400	1.129640	-0.142430	0.012789	-0.095416	0.988845	0.106571	0.031443	-0.466886	1.435443
2600	1.120310	-0.130848	0.010538	-0.080968	0.980867	0.100101	0.041620	-0.454366	1.412746
2900	1.110238	-0.116108	0.005870	-0.065165	0.976807	0.088358	0.051678	-0.436171	1.384494
3200	1.103484	-0.104130	0.000646	-0.053820	0.977982	0.075839	0.058229	-0.420177	1.361947
3500	1.099061	-0.094418	-0.004642	-0.045296	0.981948	0.063348	0.062830	-0.406573	1.343743
3900	1.095700	-0.084250	-0.011450	-0.036763	0.989471	0.047293	0.067225	-0.391662	1.324438
4300	1.089558	-0.063730	-0.025828	-0.031271	0.991000	0.040272	0.066993	-0.379206	1.312213
4700	1.087166	-0.054962	-0.032205	-0.025864	0.999188	0.026676	0.069840	-0.368687	1.298846
5100	1.086572	-0.048464	-0.038109	-0.021529	1.007795	0.013734	0.072181	-0.359986	1.287805
5600	1.087504	-0.042624	-0.044881	-0.017220	1.018597	-0.001377	0.074584	-0.351141	1.276557
6500	1.091772	-0.036881	-0.054891	-0.011667	1.036444	-0.024777	0.077745	-0.339624	1.261880
7500	1.097955	-0.032532	-0.065423	-0.007321	1.055250	-0.047930	0.080620	-0.330455	1.249836
9000	1.107385	-0.030206	-0.077178	-0.003158	1.077772	-0.074614	0.083532	-0.321832	1.238300
11000	1.118203	-0.029455	-0.088748	0.000211	1.100535	-0.100746	0.086163	-0.315018	1.228855

5.5 Exposure Compensation

From a given linear 16-bit unsigned integer wide-gamut RGB V_{sens} , the corresponding floating-point relative exposure value V_{ev} is computed by

$$V_{ev} = \left(\frac{V_{sens} - 256.0}{65535.0} \right) \frac{0.18 EI}{4.0}$$

with V_{sens} being converted to floating-point prior to black subtraction, and with EI similarly being converted to floating-point prior to its use as a scaling factor.

5.6 Logarithmic Encoding

Conversion from floating-point relative exposure value V_{ev} to normalized logarithmic encoded value V_{log} is computed by

$$V_{log} = \begin{cases} e V_{ev} + f, & V_{ev} \leq cut \\ c \log_{10}(a V_{ev} + b) + d, & V_{ev} > cut \end{cases}$$

where values for cut , a , b , c , d , e , and f are taken from the following table:

EI	<i>cut</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
160	0.005561	5.555556	0.080216	0.269036	0.381991	5.842037	0.092778
200	0.006208	5.555556	0.076621	0.266007	0.382478	5.776265	0.092782
250	0.006871	5.555556	0.072941	0.262978	0.382966	5.710494	0.092786
320	0.007622	5.555556	0.068768	0.259627	0.383508	5.637732	0.092791
400	0.008318	5.555556	0.064901	0.256598	0.383999	5.571960	0.092795
500	0.009031	5.555556	0.060939	0.253569	0.384493	5.506188	0.092800
640	0.009840	5.555556	0.056443	0.250219	0.385040	5.433426	0.092805
800	0.010591	5.555556	0.052272	0.247190	0.385537	5.367655	0.092809
1000	0.011361	5.555556	0.047996	0.244161	0.386036	5.301883	0.092814
1280	0.012235	5.555556	0.043137	0.240810	0.386590	5.229121	0.092819
1600	0.013047	5.555556	0.038625	0.237781	0.387093	5.163350	0.092824

Annex A Bibliography (Informative)

Academy of Motion Picture Arts and Sciences Science and Technology Council Procedure P-2013-001, Recommended Procedures for the Creation and Use of Digital Camera System Input Device Transforms (IDTs)

Codex DPX Multichunk User Data, Revision 15.10.2012

Color Processing in Digital Cameras, J Adams, K Parulski and K Spaulding, IEEE Micro, 1998, Vol. 18 (6)

Single-Sensor Imaging: Methods and Applications for Digital Cameras (Image Processing Series), Rastislav Lukac, Ed. (2008)

Annex B Color Rendering Example (Informative)

As mentioned in the introduction, the logarithmically encoded ARRI Wide Gamut data can be color rendered for display. There are three steps involved in this process:

- Tonemapping, in which a sigmoidal curve is applied to the logarithmically encoded ARRI Wide Gamut data, imposing a film-like ‘toe’ and ‘shoulder’, and increasing midtone contrast
- Conversion to destination primaries and chromatic adaptation to display white point
- ‘Gamma correction’, which compensates for any nonlinear electro-optical transfer function (EOTF) of the chosen display.

The following steps convert the logarithmically encoded ARRI Wide Gamut data to a form appropriate for display on a reference monitor conforming to the ITU-R BT.709 and ITU-R BT.1886 standards, or to a projector set up in accordance with SMPTE RP 431-2.

B.1 Tonemapping

A reasonable approximation to the tonemapping curve typically applied to ARRI Wide Gamut image data in the logarithmic encoding described in this document can be made by creating an interpolating spline through the following control point coordinates:

x	y	x	y	x	y
0.000	0.0000000	0.350	0.0721173	0.700	0.8137679
0.025	0.0007870	0.375	0.0970327	0.725	0.8575427
0.050	0.0016395	0.400	0.1269397	0.750	0.8932168
0.075	0.0025806	0.425	0.1619765	0.775	0.9213046
0.100	0.0036333	0.450	0.2024278	0.800	0.9425666
0.125	0.0048220	0.475	0.2484255	0.825	0.9579509
0.150	0.0063188	0.500	0.3000142	0.850	0.9685406
0.175	0.0083387	0.525	0.3571444	0.875	0.9755101
0.200	0.0110740	0.550	0.4196671	0.900	0.9800903
0.225	0.0147836	0.575	0.4873303	0.925	0.9835433
0.250	0.0198137	0.600	0.5597759	0.950	0.9871482
0.275	0.0269311	0.625	0.6337109	0.975	0.9921967
0.300	0.0375213	0.650	0.7014676	1.000	1.0000000
0.325	0.0523762	0.675	0.7616870		

B.2 Conversion for Displays with Specific Primaries, White Points and EOTFs

B.2.1 Conversion for display with ITU-R BT.709 primaries, white point and EOTF

B.2.1.1 Change of primaries

The following matrix transforms tonemapped data in the ARRI Wide Gamut color space to counterpart data for a display whose primaries and white point are those of the ITU-R BT.709 standard. The matrix is applied to the left of an RGB column vector.

Red	Green	Blue
1.485007	-0.401216	-0.083791
-0.033732	1.282887	-0.249155
0.010776	-0.122018	1.111242

This transformation is not a simple transcoding from the ARRI Wide Gamut color space to the color space of a display whose primaries and white point are those of the ITU-R BT.709 standard; it also includes a desaturation component found to provide pleasing color reproduction when applied to image data tonemapped as described in Annex B.1.

B.2.1.2 Gamma correction

The following equation should be used to transform image data V in the color space of a display conforming to the ITU-R BT.1886 standard to display-ready image data V' .

$$V' = \begin{cases} 12.00796 V, & V \leq 0.004683 \\ 1.097 V^{\frac{1}{2.725}} - 0.097, & V > 0.004683 \end{cases}$$

The pure power function of ITU-R BT.1886 (with $\gamma = 2.4$) is replaced by a piecewise function limiting the initial slope. This improves quantization when the processing is done with integer encoded values.

B.2.2 Primary conversion for display with SMPTE RP 431-2 primaries, white point and EOTF

B.2.2.1 Change of primaries

The following matrix converts the tonemapped data from the ARRI Wide Gamut color space to the color space of a projector with primaries and white point conforming to SMPTE RP 431-2. The matrix is applied to the left of an RGB column vector.

Red	Green	Blue
1.296541	-0.194182	-0.102359
0.019844	1.224098	-0.243942
0.031999	-0.036114	1.004115

This transformation is not a simple transcoding from the ARRI Wide Gamut color space to the color space of a display whose primaries and white point are those of SMPTE RP 431-2; it also includes a desaturation component found to provide pleasing color reproduction when applied to image data tonemapped as described in Annex B.1.

B.2.2.2 Gamma correction

The following equation should be used to transform image data V in the color space of a display with primaries and white point conforming to SMPTE RP 431-2 to display-ready image data V' .

$$V' = \begin{cases} 16.23957 V, & V \leq 0.003449 \\ 1.112 V^{\frac{1}{3}} - 0.112, & V > 0.003449 \end{cases}$$

The pure power function of SMPTE RP 431-2 (with $\gamma = 2.6$) is replaced by a piecewise function limiting the initial slope. This improves quantization when the processing is done with integer encoded values.

B.2.3 Primary conversion for display with SMPTE RP 431-2 primaries, ITU-R BT.709 white point and SMPTE RP 431-2 EOTF

B.2.3.1 Change of primaries

The following matrix converts the tonemapped data from the ARRI Wide Gamut color space to the color space of a projector with primaries conforming to SMPTE RP 431-2 and with white point corresponding to ITU-R BT.709. The matrix is applied to the left of an RGB column vector.

Red	Green	Blue
1.213079	-0.098707	-0.114372
0.014386	1.230503	-0.244889
0.030422	-0.021558	0.991116

This transformation is not a simple transcoding from the ARRI Wide Gamut color space to the color space of a display whose primaries are those of SMPTE RP 431-2 and whose white point is that of ITU-R BT.709; it also includes a desaturation component found to provide pleasing color reproduction when applied to image data tonemapped as described in Annex B.1.

B.2.3.2 Gamma correction

The following equation should be used to transform image data V in the color space of a display with primaries and white point conforming to SMPTE RP 431-2:2011 to display-ready image data V' .

$$V' = \begin{cases} 16.23957 V, & V \leq 0.003449 \\ 1.112 V^{\frac{1}{3}} - 0.112, & V > 0.003449 \end{cases}$$

The pure power function of SMPTE RP 431-2 (with $\gamma = 2.6$) is replaced by a piecewise function limiting the initial slope. This improves quantization when the processing is done with integer encoded values.