Wide Gamut High Performance Imaging in Digital Cinematography

Peter Centen

BU Cameras, Grass Valley Nederland B.V.
Kapittelweg 10; 4827 HG, Breda, The Netherlands
peter.centen@thomson.net
Tel: +31 76 5721564

Abstract

One of the surprises of current camera technology is the fact that at the camera output the colorimetry is still defined with respect to the displays primarily used (e.g. Rec.-709) in today's environments. An equivalent statement holds because of the mapping of the video signal into 700mV. When one disregards these standardization issues one has access to the full potential as given through the use of raw-imager-data recorded in its entirety, and then manipulated at a later point (post production) in the process. Through that, one achieves a much broader color and signal gamut than is currently being provided in today's product offerings. Leaving the signal space as defined in broadcast cameras as is, one achieves a dynamic range not found in every day use of broadcast camera. This paper discusses the above issues as a part of the VIPER FilmStream Digital Cinematography Cameras imaging ability.

1.0 Rec.-709, Wide Color Gamut RIMM/ROMM and XYZ primaries

The RGB response of a 3-imager broadcastcamera is determined by the lens, infra-red blocking filter, colorsplitter system and imagers. The signal space is definite positive, it has no negative signals. E.g. any light falling onto the imagers will in the end result in a positive output signal. After matrixing to produce colormatching to Rec-709, though, the output could go negative [1] if it where not for the blackclipper, in the camera, that limits output levels to positive numbers. This limiting operation is irreversible! It is this fact that pre-limits the color gamut of broadcast cameras.
Throughout this section a model for the optical transmission of the Red, Green and Blue channel is used to derive the {x,y}-chromaticity values for a broadcast camera. The elements in the model consist of a 2/3"-fixed focus lens, an IR-blocking filter a generic 2/3"-colorsplittersystem and 3-(generic)-imagers of which the quantum efficiency is modeled as a constant. This hypothetical camera has 3 outputs, R, G and B.
For broadcast use the RGB-output is then matched to the Color Matching Function described in Rec.-709. The matching realized through a matrix operation [1] on the raw RGB imager output signals. The (optimal) matrix M_{opt} is calculated making use of a least-square method using matrix-algebra, more elaborate models can be found in [2]:

$$M_{opt} = CMF_{709} * SEN^T * \left(SEN * SEN^T\right)^{-1}$$

With

M_{opt}	a **3x3** matrix that minimizes the least square error between
CMF_{709}	a **3xn** matrix, the wanted output signal representing the Color Matching Function according Rec.-709, and
SEN	a **3xn** matrix, representing the raw RGB signal
n	the number of test colors, which totals 80 in this simulation.

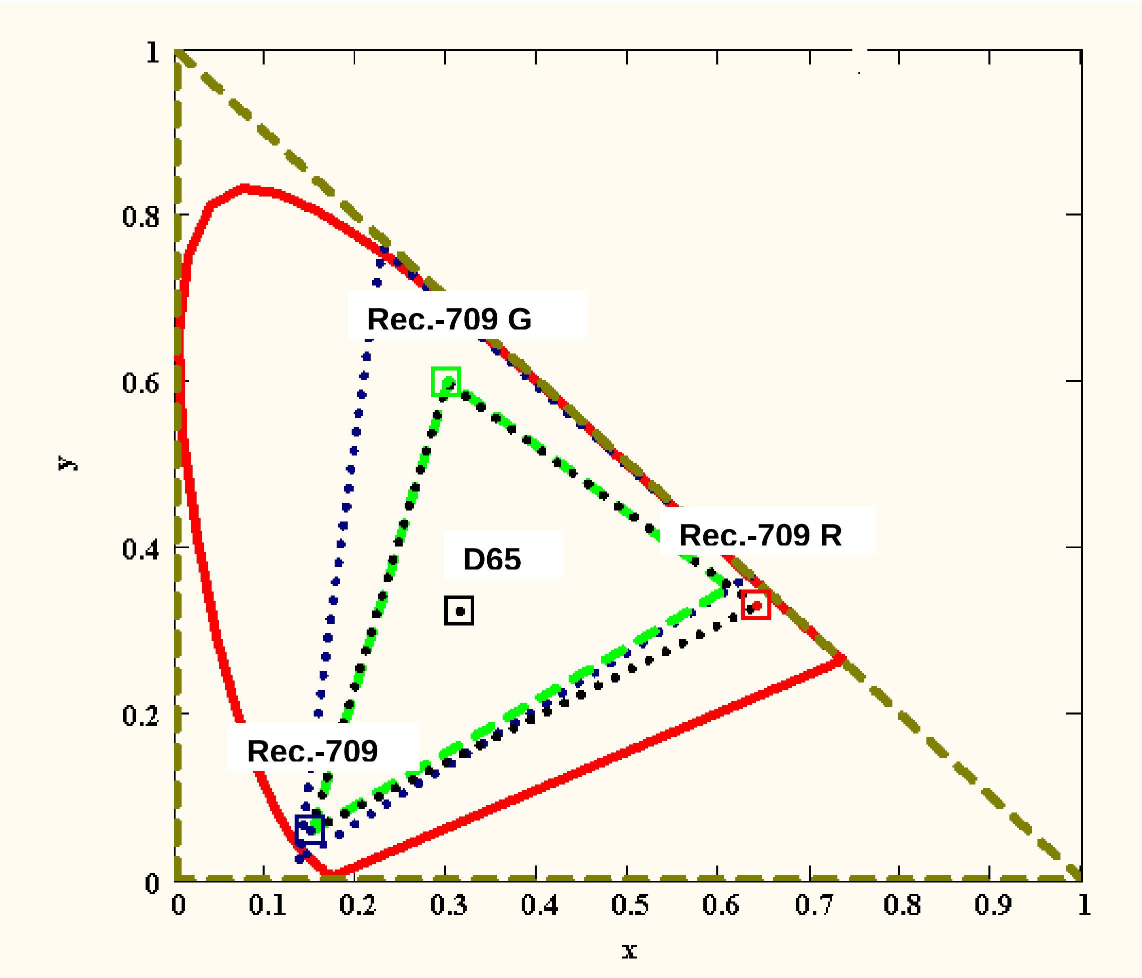

Fig 1: *The typical shape of the CIE {x,y} 2-degree 1931 observer function, red solid line, encloses the camera color gamut. The dashed (green) line triangle that nearly connects the Rec.-709 primaries is the broadcastcamera color gamut when the signals are limited to positive values only! This space can be addressed in a reversible manner even in a later -post processing- stage. The doted (blue) line which extends beyond the Rec.-709 triangle is the response one gets when also the non-physical negative video values are allowed.*

The relation between the camera-output (**RGB_{out}** a 3-element column vector) and the raw imager signal (**RGB_{sen}** a 3-element column vector) is then given by:

$$RGB_{out} = M_{opt} * RGB_{sen}$$

In this type of optimization, functions are matched as closely as possible. Or equivalently: the colorpatches used for optimization are monochromatic. Which in this case equals about (n=)80 patches. Mathematically speaking it is equivalent with having a Dirac function as input test function.

In interpreting the result one has to bear in mind that all negative video signals are mapped to zero in the RGB-channels after matrixing. With this simple camera model the color gamut spanned by a broadcast camera is calculated, figure 1. The typical shape of the CIE {x,y} 2-degree 1931 observer function, red solid line, encloses the camera color gamut. The Rec.-709 triangle is shown as the black dotted line and connects the primaries.

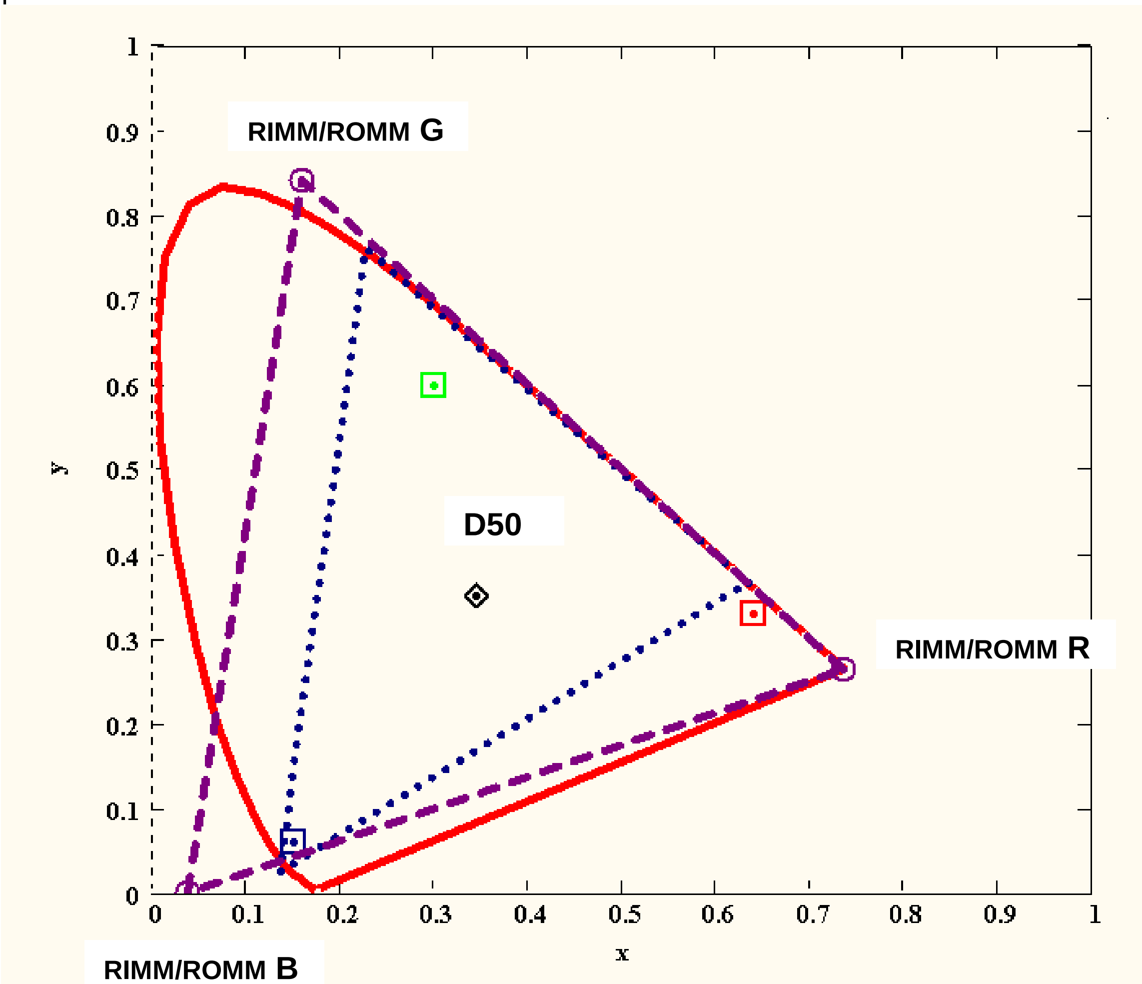

Fig 2: *The purple dashed line triangle connects the RIMM/ROMM primaries. The blue dotted line, inside this triangle and outside the Rec.-709 triangle is what is available when matched for the wide-RGB primaries.*

The dashed green line, which nearly connects the Rec.-709 triangle primaries, spans the broadcastcamera color gamut space when the signals are limited to positive values only. This space can be addressed in a reversible manner even in a later -post processing- stage. The doted blue line, which extends beyond the Rec.-709 triangle, is the response one gets when also the negative video values are allowed. Unfortunately this space is not reversible addressable in present day broadcast camera design. The non-linear operation caused by the blacklevel limiter being the cause of that. It is this more-wider-then-broadcast-color-gamut that we would like to address, too.
It is addressable in the VIPER FilmStream Digital Cinematography Camera [3].

1.1 Matching for XYZ and ICC profiling

LEMMA: The fact that all video signals are positive and matched to Rec.-709 implicitly means that every broadcast camera that exists today has a build-in rendering intent!
Let's call the raw RGB-imager data the input device color space. This color space is wider than the output color space, figure 1, the Rec.-709 triangular color space. Hence

there is a transformation between the input color space and the output color space. This is a non-linear operation, arrived at from an engineering point of view, to prevent negative video signals but with an implicit rendering intent in it. One could call it implicit color management. Since the purpose of the implicit rendering is best color matching for colors in the output color space one might call it colorimetric rendering.

There is an other way to prevent negative matrix coefficients and hence negative signal values, in the broadcast camera that exists today. Match straight against the 1931 CIE XYZ two degree observer functions. The X-camera, Y-camera and Z-camera outputs would be non-negative and could be transformed into CIEL*a*b* or CIEL*u*v* space. To be transported over the e.g. SDI-YCrCb or 2xSDI-RGB interface. Instead of the traditional RGB or YCrCb. For a (broadcast)camera the L would have the same nice properties as Y(CrCb) has, reduction of the differential gain for large signal values and nearly perceptual uniform. At the same time it would link easily to the ICC profiling discussion and would represent the profile connecting space.

1.2 Matching for RIMM/ROMM primaries

Translating the discussion about Reference Input/Output Media Metric primaries (RIMM/ROMM) for digital intermediate and showmaster into the broadcast area a much wider color gamut is possible. Application of the same optimization strategy, for deriving the optimal matrix, shows that the matrix coefficients for matching with RIMM/ROMM primaries are positive. Hence an opening is given for a potential wider gamut, the outputsignal is positive too and can be handled in a (only) positive signal space. Reversibility of matrix operation is within reach, and many other linear signal processing algorithms too. Important to mention is that one prevents the implicit rendering intent decision taken in broadcast camera matrixing.

A broadcast camera has an optical system not designed for wide gamut use. But that does not mean that it has only narrow color gamut properties, figure 2. The purple dashed line is the RIMM/ROMM primaries triangle. The blue dotted line, inside this triangle and outside the Rec.-709 primaries, is what is available when matched for the wide-RGB primaries.

2.0 The wide dynamic range and equivalent sensitometric curves of VIPER

A common used definition of dynamic range is defined as the ratio between the maximum charge an imager pixel can contain (Qmax) and the rms-noise level. A filmic approach would be to measure the average output signal as a function of exposure of the imager. When plotted on a log/log scale it would resemble the well-known sensitometric curves known from Film, an attempt to that given in [4].
The equivalent of the FOG is defined when there is no illumination applied. The associated output level is definitely not the rms-noise value and it's not zero either. The reason for this lays in the fact that in the end every camera delivers a signal that is equal or larger than zero, a definite positive signal. Let’s follow the video habit of adjusting masterblack at 0%. What actually happens is that, however small, the average black level does not equal zero but equals the average rectified noise level. Because every video processing will somewhere in the signal processing chain clip all levels below zero and hence will rectify the small noise level to an average value of

$$\frac{\sigma}{\sqrt{2*\pi}}$$

With rms-noise level equal to
σ .

Basically this enhances the effective use of a digital camera with a little over 1 f-stop. It is with this methodology in mind that the VIPER transfercurve was generated, figure 3. But this huge dynamic range is only available when the output of the VIPER FilmStream is corrected for the offset added to the RGB-signals before it was put through the 'cineon curve' [3]. Inspection of Figure 3 reveals a linear dynamic range of at least 10 f-stops and when part of the toe is put to use even several additional f-stops are available too.

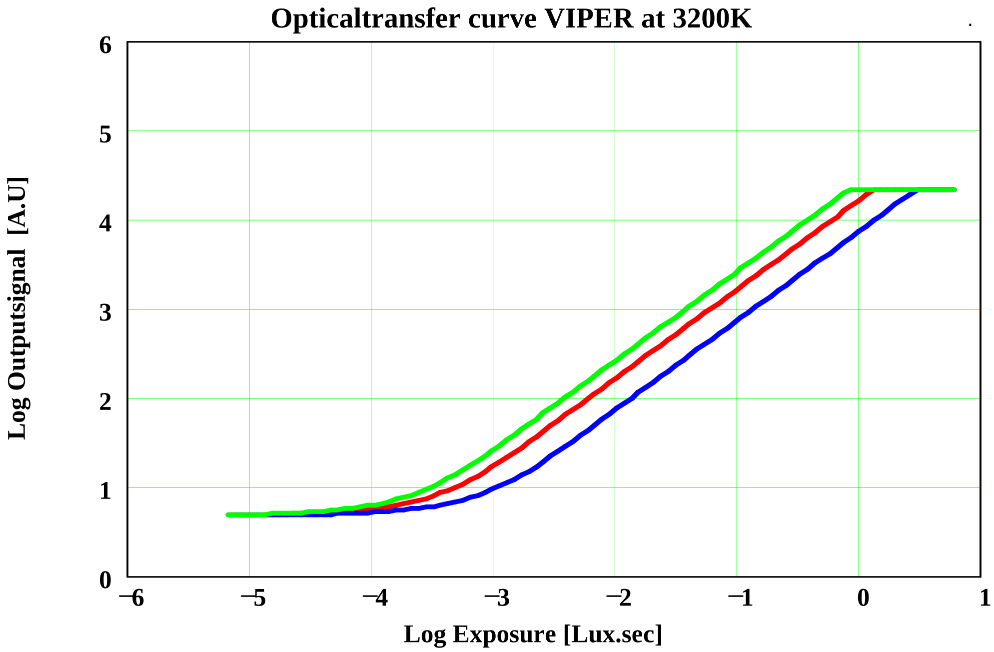

Fig 3: *The RAW data transfer curve for RGB of the VIPER. Both axes are logarithmic. The vertical axis is proportional with the log of the average output signal and the horizontal axis is the log of the exposure. At high 3200K exposures Green saturates first, then Red and finally Blue. It's this phenomenon that after white balance shows up as magenta discoloration. Notice the large exposure range over which the imager/camera generates an output signal.*

A great example of this wide signal gamut is found in Michael Mann's feature movie COLLATERAL. It has many scene's in the dark and LA-by-night as a forth character in the movie.

For clarity sake the curves in the Figures 3-4-5 are given for 3200K. But it could easily be repeated for 5000K or even 20.000K and down to 2000K it is not meant to limit. Just to clarify things.

After white balancing the RAW RGB data of the VIPER, the curves for RGB fall together for a large range of exposure values, figure 4. Depending on the type of white balance differences can be found in the low-light parts and the highlighted parts.
The left graph of figure 4 shows the electrical white balanced curves. To be executed in the camera or in post-production. The right graph of figure 4 shows optical white balancing. The vertical axis is proportional with the log of the average output signal and

the horizontal axis is proportional to the log of the exposure. In the electrical white balanced case and at high 3200K exposures Green saturates first, then Red and finally Blue. It shows up as magenta discoloration. Also the noise in blue is larger then in red and green has the lowest noise level. In the optical balanced case RGB are tracking and saturate simultaneously and give no magenta discoloration on 3200K where the electrical will. Also the noise levels will be the same for both R, G and B. The penalty though, for applying an optical white balance is about 3 f-stops in sensitivity.
Notice the large exposure range over which the imager/camera generates an output signal: at least 10 f-stops, linear and if use is made of the toe some f-stops more.

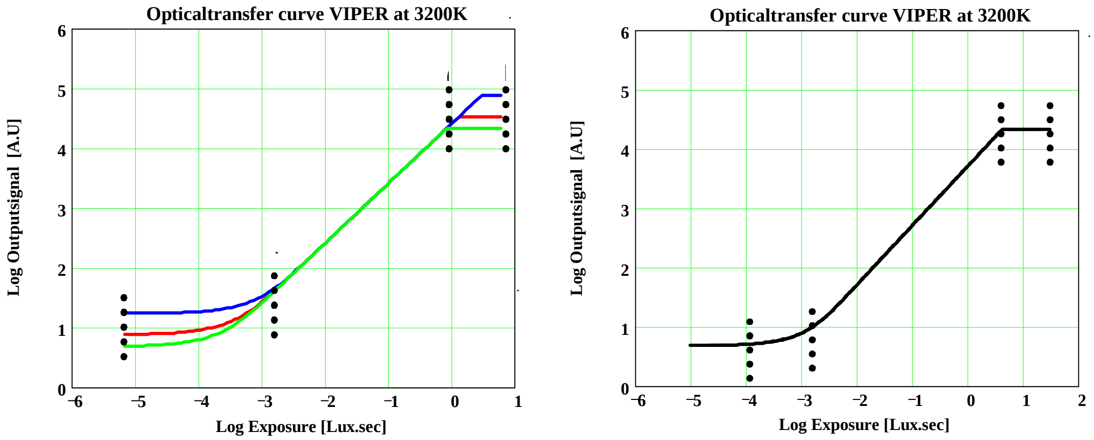

Fig 4: *The RAW-data transfer curves for RGB of the VIPER have been white balanced. Left graph shows the electrical white balanced curves. The right graph shows optical white balancing. Both X and Y axes are logarithmic. The vertical axis is proportional with the log of the average output signal and the horizontal axis is the log of the exposure. In the electrical white-balanced case (left graph) and at high 3200K exposures Green saturates first, then Red and finally Blue. It shows up as magenta discoloration. Also the noise in blue is larger then in red and green has the lowest noise level after electrical white balancing. In the optical balanced case (right graph) RGB are tracking and saturate simultaneously and give no magenta discoloration. Also the noise levels will be the same for both R, G and B. The penalty though for doing an optical white balance is about 3 f-stops.*

2.1 To loose and gain at the same time: Dynamic range in a broadcast camera
LEMMA: The introduction of the gain switch in a camera improved operationality but reduced signal gamut, known as dynamic range.
A camera used in 0dB will have a signal range as depicted in the upper graph of figure 5. The left lower graph shows the use in +12dB and +24dB usage is shown in the right lower graph.
The upper graph of figure 5: At low illumination levels the output is a mix of an offset, caused by rectified noise, and the light generated image signal. When exposure increases the output voltage increases too in relation to the amount of light falling onto the imager. It is getting increasingly linear with exposure. The other end of the curve is where saturation occurs. This saturation in an imager is known as Qmax. In contrast with Film an imager does not saturate gracefully.
At 3200K green starts to saturate first then red and finally after some 2 f-stops more exposure Blue. The low blue signal content a consequence of the 3200K light. The nominal set point is defined as the 0dB sensitivity calibration.

All output levels above the nominal set point are in a region, which is at best described as compression region, known as the knee region [4,5]. In figure 5 the range encompasses 3 f-stops.
Switching the camera to +12dB gain results in a decrease of signal gamut. One limits it! Since only a part of the total available response curve is used now, lower left graph of Figure 5. The dynamic range is limited with 12dB or 2-fstop. This is even more pronounced for +24 dB gain. Now the dynamic range is even more limited with an additional 4 f-stops. This holds for every broadcast camera!
Of course the operational demands for this are clear. But why should one want to use it in a camera like a VIPER.

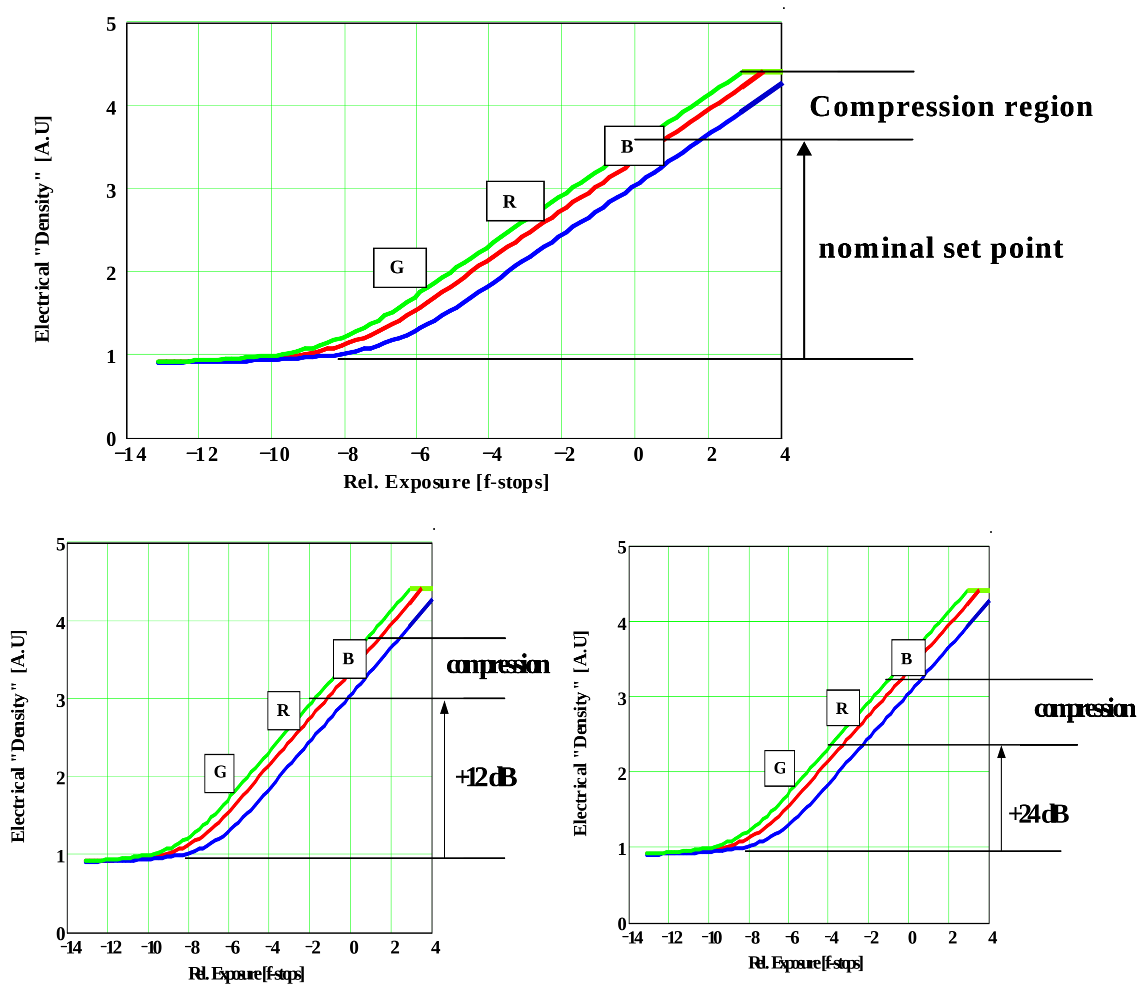

Fig 5: *Three graphs showing the output ranges that are available when the camera is switched to 0dB (upper graph), +12dB in the lower left graph and +24dB in the lower right graph. The compression region has a range of 3 f-stops above the nominal operating point. Note the fact that at higher gains a smaller part of the total response curve is used!*

2.2 How many bits are enough?

LEMMA: In an imager-pixel electrons are generated and the number is proportional to the amount of light falling on that pixel. Since there are only 'whole' electrons the pixel-output signal is quantized in steps of 1-electron. So the maximum number of bits needed to represent a signal from the imager may always be limited to $^{2}\log(Q_{max})$.

Where Qmax is the maximum number of electrons a pixel can contain.
Many discussions in the broadcast industry and in the Electronic Cinema are devoted to the number of bits needed for representing the video signal. In this paper the approach taken is based on first principles and derives an upper limit.
Every photon that is absorbed in the silicon will generate an electron-hole pair. Depending on collection depth and effective pixel-aperture [6], usually, a large fraction of the generated electrons are collected in the pixel. The remaining ones are drained to the substrate.
The collected electrons in the imager-pixel are called a charge packet. When the pixel is readout the charge in converted into a voltage through the use of a capacitor. The conversion gain, which is a well-known parameter when specifying imagers, is known as the number of **μV/e**. Presently one finds numbers between 10-50 μV/e.
Assuming 20 μV/e then every electron more in a charge packet results in 20 μV more output signal. The signal is quantized in steps of 1-electron and in this example in steps of 20 μV. It is only the thermal noise, which is added afterwards, that disguises this quantizing nature.
Even the shotnoise in the charge packet is quantized in steps of 1-electrons.
An example: A 5x5μm^2 square pixel will have about 4000 electrons generated in the green-pixel when viewing a scene illuminated at 2000, f/11, 3200K and having 90% reflection. Assuming further a maximum of 5 times overexposure one arrives at a maximum of 20k electrons. As stated before the charge packet is quantized in steps of 1-electron, so the signal is quantized too. Hence the maximum number of levels in this example is 20.000 and can therefore always be represented with a 15bit word!
Remember that in this example the effect of the thermal readout noise, which is always added to the signal, is not even taken into account. The effect of noise will always be to reduce the number of bits needed to convert the signal to digital numbers [3]. The basic question to be answered is then: why should one want to use the number of bits to enhance the accuracy for expressing the noise.

3.0 Wide gamut MTF in relation to Imager-PIXELS and Film-GRAIN

LEMMA: The imager-pixel does not change shape when exposure changes hence its MTF in RGB will be independent of exposure. The large film grains are tagged at low exposures and the smaller one's at higher exposures. Hence the MTF of Film must depend on exposure.
The above provocative statement also answers the question of why changing the gain in a camera effectively changes the ASA number. Film stock only has one ASA number that does not change when applying gain in post.
It also includes that the granularity of the noise in an imager-equipped camera is not different in the low-exposed part compared with the higher exposed ones.

The light sensitive element in an imager, the pixel, has a fixed geometry. All pixels are equal in every frame and independent of exposure. Together with a linear conversion process between number of photons and electrical signal the imager is a very wide dynamic range linear light to electrical signal converter. With film the tagging of the grains is a statistical process but on average the larger grains are tagged at low exposures and the smaller ones at high exposures. From a MTF point of view a large grain resembles a large aperture and a small grain a small aperture. From imagers it is known that the aperture determines both sensitivity and MTF, hence it must have some effect.

Conclusion

It was argued that when sticking to the broadcast camera with its matrix to conform to the Rec.-709 colorimetry the color gamut is needlessly limited. It finds its cause in the fact that negative signals are clipped in the camera. Hence there is no undo possible of those parts that are clipped already in the camera. At the same time that fact includes an implicit rendering intent.

The colorimetry of the VIPER Film Stream Digital Cinematography Camera is extended beyond the Rec.-709 chromaticity triangle. It is accomplished by just taking the raw imager RGB-output and do the white balance and chromaticity mapping in post.
One could even define a matrix that is optimized for the RIMM/ROMM wide gamut RGB-primaries or for XYZ-primaries. The resulting matrix coefficients are positive and one has access to the full color gamut.

At the same time it is acknowledged that in the present broadcast studio environment it is needed to have some prescribed mapping for the colorimetry in order to ease the burden on everyday working practice.

In high gain use of a broadcast camera the effective dynamic range is reduced with the same amount the gain is increased. The usable signal space is reduced with that same amount.

Due to the property that an image-pixel has a fixed geometry regardless of exposure levels it can be used from very low light levels onwards.

Discussed was the fact that an imager pixel is an electron-counter and as such quantizes the signal in steps of 1-electron. This was, in an example, related to the number of bits needed to represent imager signals and an equation was given to calculate the upper bound to the number of bits.

LITERATURE

[1] C. Poynton;*Digital Video and HDTV Algorithms and Interfaces*; Morgan Kaufmann Publishers 2003; ISBN 1-55860-792-7.

[2] G.D. Finlayson, M.S.Drew*; White-point preserving color correction;* The fifth Color Imaging Conference, 1997, pp:258 – 261.

[3] J.v. Rooy, P. Centen, M. Stekelenburg,*The VIPER Filmstream Camera: A marriage between Film and Video,* proceedings of the Broadcast Asia Technical Conference and Exhibition, Singapore, June 18-22, 2002.

[4] L.J. Thorpe; *HDTV and Film – Digitization and Extended Dynamic Range* ; SMPTE journal, June 1993, pp 486-497

[5] J. van Rooy, R. Voet, *Twelve-Bit Acquisition, The Next Step in Digital Broadcast Cameras,* SMPTE journal, October 1997, pp 698-704

[6] A.Theuwissen. *Solid-State Imaging with Charge-Coupled Devices*. Kluwer Academic, Publishers 1995; ISBN 0-7923-3456-6.