Viper FilmStream Camera: A Technical Overview

Jan van Rooy, Peter Centen, Mike Stekelenburg

Abstract

This paper proposes a camera for a new workflow in which picture data from the CCDs of the camera are transferred directly into postproduction, maintaining the best possible creative freedom for post-production. The picture data in this FilmStream workflow are transferred to postproduction using standard interfaces, a logarithmic transfer curve and 4:4:4 sampled RGB data.

FilmStream uses open standards like dual HD-SDI (SMPTE372M), with fixed gains in R,G and B and a predefined logarithmic correction to convert the 12 bit linear CCD signals to a 10 bit logarithmic signal. The paper will describe the camera, the workflow and the interfaces necessary to achieve an entirely new way of production[1].

Introduction

The main purpose of any imaging system is to give a convincing representation of reality on a standardised display. However, displays cannot cope with the large dynamic range existing in many scenes. Convincing means that the imaging system has to provide tools to render a picture according to the intent of the director, taking into account the inherent limits of the display that will be used.

Traditionally the workflows of television and film production have been very different. Video production aimed at achieving the best possible picture for the targeted viewing device immediately in the camera, while the film approach separates distinct stages of first capturing an image with a wide dynamic range and then manipulating this picture in either filmlab or post-production to serve the more limited range of the film screen.

Borders between both ways of production have blurred. Electronic capturing of images has found its way to feature film production, and electronic post-production is used in film originated production. In the video world pictures see more and more post-production operations before final transmission.

This paper describes a camera for a new workflow in which picture data from the CCDs of the camera are transferred directly into postproduction, maintaining the best possible creative freedom. The picture data are transferred using standard interfaces, and 4:4:4 sampled RGB data. This interface is called FilmStream and will be described in more detail.

… a marriage between Film and Video

FILMSTREAM™

THOMSON MULTI MEDIA BROADCAST SOLUTIONS

General Camera Architecture

The Viper FilmStream camera uses a 3 CCD concept, with the FT25 multi format sensors. It can deliver pictures according to several formats, and two aspect ratios (16:9 and 2.37:1). The use of a 2/3-inch format guarantees a compact camera design, and the availability of a wide range of lenses and accessories. Table 1 gives a summary of the main properties of the camera,

	Viper FilmStream Camera
On board recorder	no
Imaging Modes	FilmStream, RGB, YUV
Luminance resolution (H)	1920
Chroma resolution (H)	FilmStream & RGB: 1920 YUV: 960
Colour primaries & white point	FilmStream: Uncommitted RGB & YUV: ITU Rec. 709
Quantization	10 bit
Compression	ZERO
Format	1080 active lines (720 also available)
Aspect ratio	16:9 & 2.37:1
Frames per second	23.98, 24, 25, 29.97, 30, 50, 59.94, 60

Table 1: main properties of the camera. A switchable solution has been designed that serves the two modes of operation. Figure 1 depicts these two possible modes of operation from a conceptual point of view. The thick arrow represents the user interaction for picture control

The Filmstream mode

In the FilmStream mode the signals from the CCDs are converted to 12 bit digital RGB signals. The full dynamic range of the CCDs is output for recording. The viewing channel processes these wide gamut signals the same as a video processing channel in a broadcast camera would do, but here it is only used as a monitoring function. It gives an impression of what the end result might be after basic image processing. This impression can be used by the director of photography to tweak his lighting, and note the basic settings of the viewing channel as a guideline for postproduction. If viewing at the set is considered unnecessary, the DoP can decide not to use the monitoring function at all and simply use the lightmeter to adjust the exposure of the camera. Due to the linear-in-light nature of the output it is possible to use the camera at various exposure indexes, but the range of approximately 300-400 ASA gives the best optimum between highlight range and noise

As the preview will be on a monitor we choose to use the standard camera controls for the viewing channel. White balance, camera gain, knee, gamma, matrix, contours and several other functions are set at the camera, and the end result of these processing stages is viewed at the set. The target display is the standardized CRT screen with rec709 colorimetry.

The camera's viewing channel acts as a dynamic range compressor to cope with the limitations of the CRT. Actually, any post processing in the main post-production chain does essentially the same thing: adapting to the smaller dynamic range of the display device.

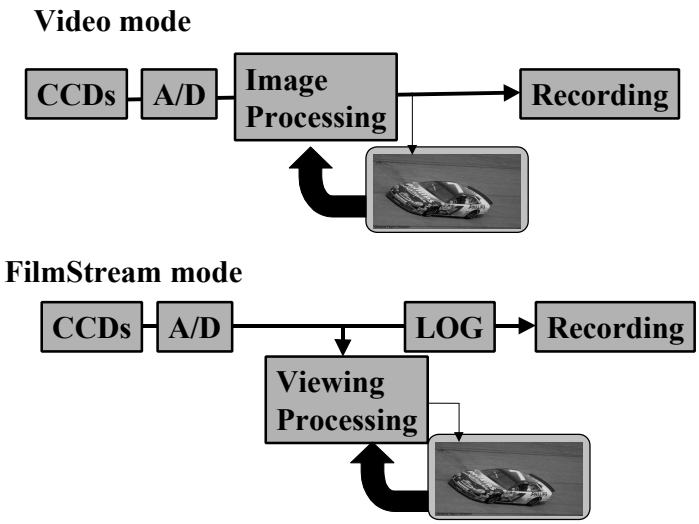

Figure 1: Video and Filmstream modes of processing

Description of the video mode

In the video mode the signals from the CCDs are digitised in the same way as in the FilmStream mode. These digital signals are processed and viewed on a monitor. But here the processed pictures are output to the recorder. The camera is a full-featured HDTV camera, which is capable of outputting many standards[2] (see table 1). The output is either YCrCb on a single interface, or RGB 4:4:4 on dual link HD-SDI. The video mode is a perfect mode of operation if you like to decide the final looks of the scene at the set, as postproduction will not be necessary. The Viper FilmStream camera provides a full range of powerful menu driven tools for this mode of operation[3]. The image processing controls are adjusted until a satisfactory result is viewed on the monitor and the end result is recorded and used. It is only possible to adjust before recording. The dynamic range is adjusted for the limited range of the target viewing device and typically will be lower than the full range of the CCDs.

The 4:4:4 RGB mode has full spatial resolution and is ideal for critical post-processing like color keying and insertion of graphics. However, creative freedom after shooting the scene is limited with respect to dynamic range and color correction. When choosing for this video way of operation and 4:2:2 YcrCb sampling, tape based recording can be used. For instance uncompressed YCrCb on the Voodoo recorder or compressed on HD-D5.

The Filmstream Interface

Physical interface

To integrate smoothly into the present tools for postproduction the interface from the camera to the postproduction world has to match existing standards as closely as possible. The data rate of the signal from the camera is too high for a standard HD-SDI link, so a dual link approach according to SMPTE 372M has been chosen to transport the 3x10 bit RGB data. It is possible to transmit 3x12 bit within the standard, but 3x 10 bit fits better in the 32bit storage structure used in most computers. Most high-end postproduction systems can handle the dual link signals from the camera or recording device.

Transfer curve

It is desirable to have a 3 x 10 bit interface for recording and postproduction, simply because 3x10 bit is less than 32 bit and 3x12 bit is not. This implies that we must find a way to transform 12 bit linear signals into 10 bit signals in a transparent way. The noise characteristic of image sensors suggests that coarser quantisation is allowed in the upper range of the sensor signal. See appendices 1 and 2 for a short explanation of sensor and quantisation noise.

The noise contribution of the sensor and the noise of the log quantisation are drawn in figure 2.

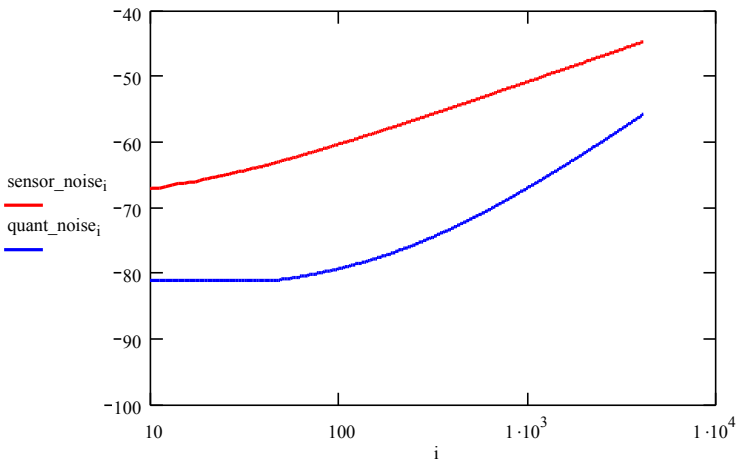

Figuur 1: sensor noise and quantisation noise

As can be seen in figure 2 the quantisation noise of the log curve is constantly more than 12 dB below the sensor noise.

This means that mapping 12 bit into a 10 bit log curve has a noise contribution of about 0.25dB or less, and can be considered as visually loss-less, with the quantisation noise fully de-correlated from the signal.

Although the optimal matched curve to the noise characteristics of the sensor would be something like a modified square root function, there are several reasons to prefer a log curve:

1. Working with logarithmic curves is an established practice in postproduction[4]
2. SMPTE is in the process of standardising log representation

Although a log curves approach transforms the linear light signal from the sensor into a signal that is linear in perception, that is not the main objective. It may be convenient in post-processing, but transferring the 12 bit linear light signal to postproduction through a 10 bit interface & storage device is the most important part.

To match the chosen curve to the specific properties of a sensor signal a number of choices have been made to allow for headroom. In order to avoid clipping in the signal, the absolute black level and peak white level in the linear sensor signal are at 64 and 3840. This allows some room for tolerances and noise in the signal. The number range of the chosen log curve excludes the forbidden values of the serial digital standard.
The specification of the transfer curve is as follows:

Camera log curve

output of CCD + A/D converter:
    linear out, 12 bit 0..4095, black level set at 64, sensor max level set at 3840
log curve and B/W clip
    log curve (=Cineon Curve):
        if x> 37 then y= 500*log(0.02714*x) else y=0:
    black and white clipper:
        if y<3 then y=3
        if y>1020 then y=1020

Inverse curve

In post processing the FilmStream can be transferred back to a linear-in-light signal using the inverse transform:

$$x=10(y/500)/0.02714$$

The result is a linear representation of the sensor signal with black at 64 in a 12 bit range.

It should be noted that that this result is fundamentally different from the result that negative film gives when scanned into postproduction. The gamma of negative film is approximately 0.64, and thus does not have a linear response. The linear-in light signal of FilmStream permits operations like white balance, primary color correction, camera gain to be performed in an easier and mathematically more correct way.

Recording

The FilmStream method generates a lot of data from the camera. In order to capture the full dynamic range of the sensors, and to connect to postproduction the following is needed:

- 12 bit A/D conversion at the sensor
- full 1920x1080 4:4:4 sampled RGB signals
- 12 bit linear to 10 bit log conversion
- A dual link RGB 4:4:4 interface

Recording can be done with disk arrays of several manufacturers, there are also several capture cards available (SGI, DVS) capable of accepting a dual link HD-SDI signal. There is at least one portable recording solution on the market that can record FilmStream signals[5]

The postproduction tools from most leading manufacturers are able to work with the 4:4:4 RGB data coming out of the camera.

Filmstream: Practical Issues

FilmStream in a 2K post-production world

Comparing the route from the scene to postproduction with the Viper compared to film, the results for 2K resolution in image sharpness should be alike.

In film there are two optical processes, one in the camera and another one in the film scanner. The negative film has a gamma of 0.65. The image is sampled with a horizontal resolution of 2K pixels, and proper optical measures have to be taken to avoid aliasing.

The Filmstream workflow has only one optical process. The camera samples the picture with a horizontal resolution of 2k pixels, and a comparable optical filtering has to be employed as in the 2k film scanner. Light is translated into a signal level in a linear way. A linear representation has advantages in further processing, for instance white balance, black balance and resampling and filtering the image data[4].

Dynamic range

The dynamic range of FilmStream is determined by the dynamic range of the sensors. Present state of the art is that HDTV sensors have a signal to noise ratio of approximately 54dB and a headroom of about 2 f-stops This gives a theoretical dynamic range of 66dB, or more than 10 f-stops.

Conclusions

We introduced a new camera concept and a new workflow called FilmStream that enables us to record the image data directly from the CCDs. This maintains the maximum creative freedom for the cinematographer to get the highest quality pictures. It is possible to experiment with different settings, even after the material has been recorded.

The camera also allows a workflow where many decisions are made directly at the set, and the processed picture is recorded in the video mode.

In both modes full spatial resolution is maintained, and resolution will be comparable to film scanned at 2K resolution.

Appendix 1: CCD Noise and Exposure[6]

Electronic image sensors integrate the photo current of a photodiode during a certain exposure time. The electrons generated are converted to a voltage on a capacitance. The integration capacitance is constant, so the output voltage is typically linear with light over the output range.

In a CCD imager, sensor amplifier noise and signal shot noise are the dominant noise sources. In black, the sensor amplifier noise is the dominant factor, with increasing light levels on the sensor the signal shot noise becomes the dominant factor. It is common in image sensor to express values in electrons.

Sensor amplifier noise

The sensor amplifier noise n_black is generated by the on chip amplifiers and is not dependent on the exposure of the sensor. Typical values for a well designed CCD is about 10-15 electrons rms.

Signal shot noise

Shot noise is proportional to the square root of the number of electrons generated in a pixel.

The dynamic range of a HDTV sensor capable of delivering 54dB signal to noise ratio and 600% overexposure is 54+15= 69dB

Or in electrons: the maximum number of electrons N_max is n_black + 69 dB

In the following figure the noise (in electrons RMS) as a function of exposure (in electrons) is plotted

n_black=10 el and N_max=30 kel

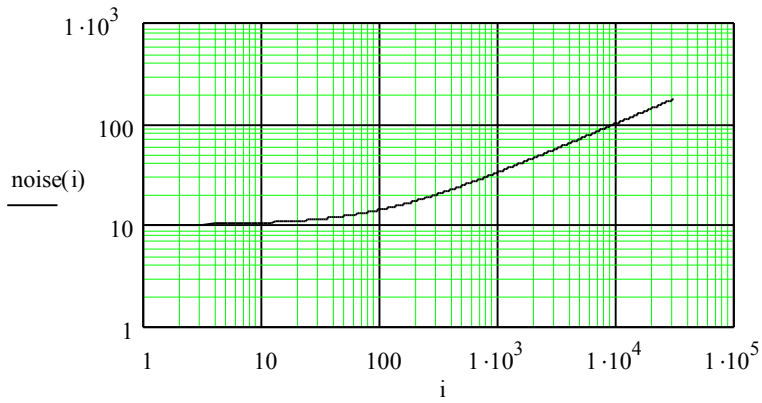

Figuur 3: noise as a function of exposure (both in electrons)

Appendix 2: Quantisation and Noise

In quantising a signal with an A/D converter, we produce errors with respect to the original analog signal. These errors are known as quantisation noise. If we use a quantiser with N levels, the signal to noise ratio of the quantising process can be described as[7]:

$$S/N = 20 * \log(N) + 10.8 \text{ dB}$$

Or if we generalise this to the relative stepsize of one quantisation step versus the total range Spp:

$$S/N = 20 * \log(Spp/stepsize) + 10.8 \text{ dB}$$

Here the signal S is expressed as a peak to peak signal, and the noise N as rms noise.

If the rms noise level in the input signal is more than 0.3 quantisation levels the quantisation noise can be considered to be uncorrellated to the signal and can be simply treated as an independent noise source. A model of A/D conversion can be found in reference[8].

1 More information can be found on www.viperfilmstreamcamera.com
2 Centen et al, 2001. A multi format HDTV camera head SMPTE journal august 2001 pp 510 to 516
3 Jan van Rooy and Rob Voet, 1997 Twelve bit Acquisition, the next step in digital broadcast cameras SMPTE journal october 1997 pp 698 to 704
4 Glen Kennel an David Snider, 1993. Gray-scale transformations of digital film data for display, conversion and film recording SMPTE journal december 1993 pp 1109 to 1119
5 HDREEL leaflet, see www.directorsfriend.de
6 B.Dierickx, Electronic image sensors versus Film: beyond state-of –the art. 1999 http://phot.epfl.ch/workshop/wks99/2_1.htm
7 C.P Sandbank (ed) Digital Television 1990 pp 27 to 30
8 L.Rabiner and B.Gold, Theory and application of digital signal processing, 1975, Prentice Hall, ISBN 0-13-914101-4 section 5.2