

White Paper

What Digital Resolution is Needed to Scan Motion Picture Film: 4K, or Higher?

In the last few years, both digital intermediate (DI) postproduction systems and digital projection have advanced to full 4K resolution, avoiding the generation losses of traditional photo-chemical workflows. This leads to the question of how high a resolution is necessary to scan the film at the beginning of such a production chain.

In fact, with the possibility that further improvements in both the DI and projection stages even beyond 4K could occur in future, it seems reasonable to conclude that the only limitation that should be applied at the scanning stage is the information-carrying capacity of the film itself.

This paper examines the questions of just what that limit is, what the required parameter values are – both analog and digital – to capture it, and what practical issues are involved in designing a film scanner that pursues such values.

The Information Capacity of 35 mm Motion Picture Film

Although larger optical formats can capture more spatial information, the predominance of 35 mm in motion picture cinematography leads us to the assumption that a film scanner should primarily consider the image frame formats associated with this gauge.

35 mm film can hold an extremely high density of information. In calculating how much in digital terms, it is necessary first to begin with the Modulation Transfer Function (MTF) characteristics of a particular sample. This complex analog quantity must then be transformed into a digital equivalent, via the application of sampling theory.

MTF on film

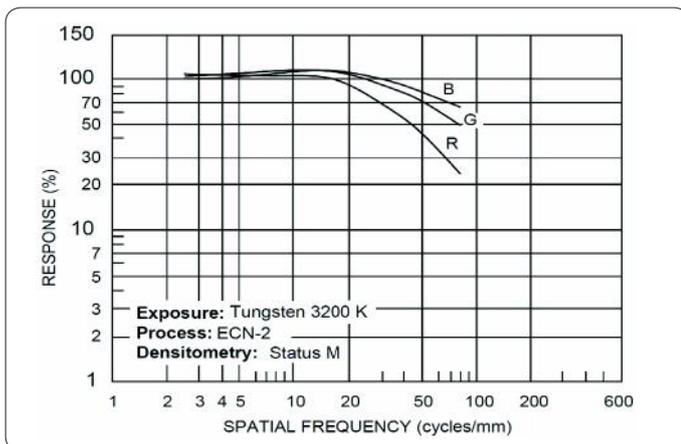


Figure 1 – MTF on Camera Negative Film (Eastman Kodak Vision 3 5207 250D)¹

First, we can consider the published data for the MTF of a particular film stock. We can assume we will find the highest readings on a camera negative. The example shown in Figure 1 is taken from the manufacturer’s data for Eastman Kodak 5207 250D color camera negative. Modulation in the three color layers is plotted out to 80 line pairs/millimeter (lp/mm) at a level falling to just under 50% in the green recording layer, with blue a bit higher and red somewhat lower.

However, this is not necessarily the level of modulation that would be obtained in practice. To expose an image on film (other than film-recording it from a computer file), we have to get it there via a camera and lens. Lenses have MTF responses, too; the image resulting on the film is therefore a convolution of the lens characteristics with the film’s own response. A useful test that incorporated this fact was conducted by the ITU in 2001-2002 as part of a project known as “Large Scale Digital Imagery”. Although the overall objective was to examine film answer print and release print resolution, the data that

was collected also included measurements of the MTF of the original camera negative (OCN) (Figure 2). In this test, the ITU measured a close-to-limiting* modulation depth of 6% at 106 lp/mm, on the OCN (Eastman Kodak 5274). The ITU’s report also included information on the camera lenses used. Although these were normal production lenses, they were set to a fixed aperture for optimum MTF, i.e. minimizing resolution losses from both aberrations and diffraction (a constraint that might not always be possible in a normal production).

* Compared to ISO 12233 definition.

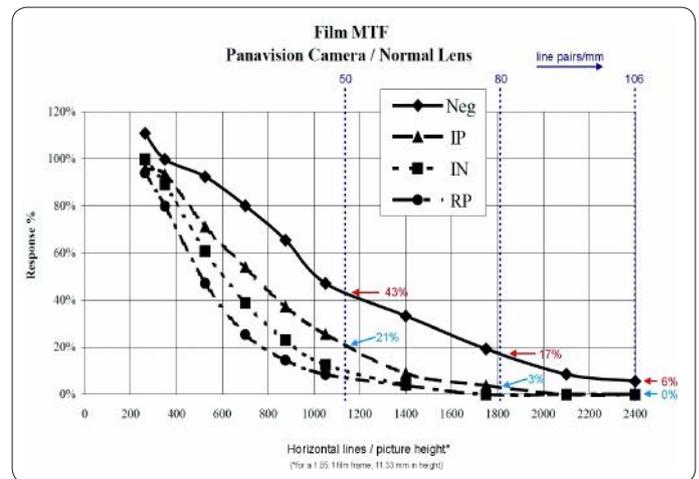


Figure 2 – MTF on Camera Negative Film (figure based on ITU test results)²

Notes:

1. The „horizontal lines / picture height“ scale is format-sensitive, here based on a frame 20.96 mm x 11.33 mm (1.85:1). Slightly higher numbers would occur with a Super 35 mm frame, but the added line pairs/mm scale is absolute and can be compared with the scale in Figure 1.
2. The „Normal lens“ means spherical, as opposed to an anamorphic lens used in another part of the test sequence.

It could therefore appear that rather than the highest spatial resolution indicated on the manufacturer’s data sheet of a current film stock (80 lp/mm at 50%), the close-to-limiting resolution of 106 lp/mm at 6% found by the ITU might be a possible target for capturing “all” of the information on the highest quality 35 mm film. However, this has to be considered in the context of three inter-related parameters: limiting resolution, sharpness, and aliasing, because these are the factors that concern us when we make the transition from the analog information on the film to a digital representation in the scanner.

Limiting Resolution, Sharpness and Aliasing

At an MTF of 6%, the ITU’s measured result at 106 lp/mm on the developed OCN was evidently close to the theoretical limiting resolution of the film stock; the response at much lower spatial frequencies – 20 to 50 lp/mm – has a better correlation with perceived sharpness.

But what is important about the limiting resolution is the potential for any modulation at this frequency to induce visible aliasing when scanned with a digital sensor. Avoiding such aliasing is the most important factor in deciding the necessary digital resolution of the scanner.

Nevertheless, the conclusion has to be that reading film information at 106 lp/mm is not of the highest significance, because it is at too low a level either to be visually significant or to trigger visible aliasing. Instead, the ITU's measurement at 80 lp/mm in the same curve seems more meaningful, because it is higher in level at about 17%*, and also because it confirms the validity of the limit of 80 lp/mm plotted on the Kodak curves in Figure 1.

Limiting the target resolution to 80 lp/mm rather than 106 lp/mm, therefore, makes sense, because the reduced MTF via the camera lens has already led to a much lower modulation level, so that any aliasing that does occur from frequencies beyond this will effectively be invisible (provided the scanning sensor's pixel layout is appropriately chosen - see below).

From MTF to Scanner Resolution

We need to find the number of pixels required on the scanner's sensors to read 80 lp/mm on the film and describe the information in the popular "K" notation**, i.e. quoting only the horizontal axis.

We must therefore consider how many millimeters' total distance we are scanning horizontally. For Super 35 mm format, this is 24.92 mm across the exposed frame width, meaning that we need enough horizontal pixels to read 80 x 24.92 or 1994 line pairs total. Sampling theory tells us that since a "line pair" is one complete cycle of a sine wave, Nyquist frequency for the sensor will equal the line pair count per scan line, and therefore pixel frequency will be a minimum of twice this, or 3988 pixels per scan line. Is our answer therefore that we need at least a "4K" scanner?

Before we conclude that "4K" is indeed the answer, let's look a little more closely at the scanning function.

Another MTF to Consider!

A digital scanner is "analog" in one sense: its sensor has its own MTF. This arises because each digital sample has to be created by looking through a "window" at the continuous information on the film. The "window" is of course an individual pixel, but because the pixel is required to measure just one level for the whole of its window, it must average all the variations occurring within the window. The window therefore has its own MTF, which is convolved with the MTF of the information on the film (which is itself a convolution, as discussed earlier), reducing the detected level of the detail on the film even further. In the context of an image sensor, this function is also known as the "geometric MTF" of the sensor (to distinguish it from other sources of resolution loss in solid state image sensors). Like the MTF curves for the film, the geometric MTF curve has a limiting frequency and a shape. Unlike the film MTF, however, the geometric MTF curve has a very regular shape and a very clearly defined limiting frequency. Because sampling is involved, the geometric MTF curve also has an alias curve associated with it, also of very regular shape and extent. However, the layout of the pixels has a profound effect on the geometric MTF curve and its aliasing, as will be seen later.

And Yet Another MTF!

Between the film and the sensor is the scanner's own lens, which has its own MTF, too. However, since this lens operates under completely fixed geometry, with a magnification factor close to unity, and with very favourable lighting conditions, it can avoid the optical compromises inherent in most camera lenses.

For example, it can be set to operate with a fairly small aperture, thus making any lens aberrations insignificantly small, while the relatively large optical format of the film and sensors means that diffraction losses are also very small. Furthermore, defocusing loss with irregular film can be minimized via a large depth of field. In total, therefore, the convolution of the scanner lens MTF with the other MTFs can be designed to be quite insignificant (see Figure 5).

* If the ITU's result at 80 lp/mm of 17% is compared with Kodak's published figure of about 50% for the same frequency, the difference may seem large. However the MTF difference can be explained via two factors: first, camera film stocks have advanced in performance in the years since the ITU test (limiting resolution remains similar, but modulation depth at given spatial frequencies has increased appreciably; to see this, the published data for the 5207 camera stock shown in Figure 1 should be compared with that of the 5274 stock used in the ITU's test). The second factor is the way the film was exposed (through a good-quality practical production lens, through a scientific diffraction-limited lens, or using no lens at all?). In any case, the two sets of data merely illustrate possible targets for a film scanner to aim for, and should not necessarily be compared directly.

** Example: "4K" means 4096 pixels of horizontal resolution, "2K" means 2048 horizontal pixels, etc. Vertical pixel count is not stated, because it can be calculated from the aspect ratio, since pixels in film scanning are usually square.

Scanner Pixel Arrangements, MTF and Aliasing

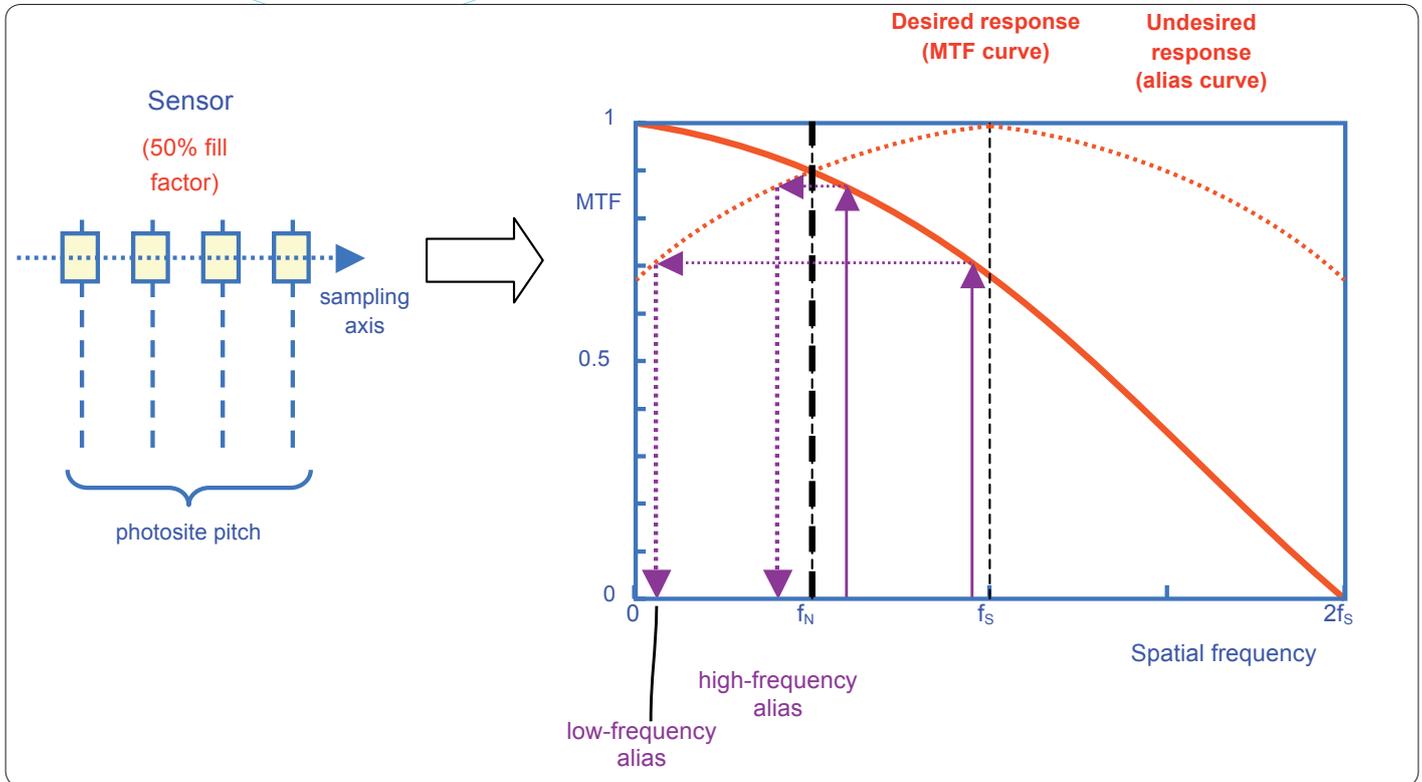


Figure 3 – Sensor with 50% Fill Factor

The left side of Figure 3 shows some pixels (photosites) in a scanner sensor, and on the right the resulting geometric MTF and alias responses with no prefiltering, i.e. input frequencies are allowed to extend beyond the Nyquist limit f_N to sampling frequency f_s and beyond. The Nyquist limit is a function of the pixel pitch: the smaller the pitch, the higher the Nyquist frequency. The geometric MTF (solid curve) of this layout is quite high (90%) at f_N , but the undesired alias is also 90% at f_N and does not decay very rapidly back towards zero frequency. This is a consequence of the particular sensor layout, where the shape of both MTF and alias curves is governed by the photosite fill factor³, in this case 50%.

Considering the effect of the 50% fill factor on aliasing:

- a signal frequency below f_N will theoretically produce no alias
- a signal frequency not far above f_N will „wrap around“ as shown to produce a high frequency alias
- a much higher signal frequency – close to f_s – will produce a much lower-frequency alias.

The concern here is that from f_N onwards the alias amplitude is the same as that of the signal frequency that causes it. While the 50% fill factor layout gives a high geometric MTF, it also produces high alias amplitudes. Most seriously, the amplitude remains high in alias frequencies close to zero, which are much more visible than aliases at high frequencies.

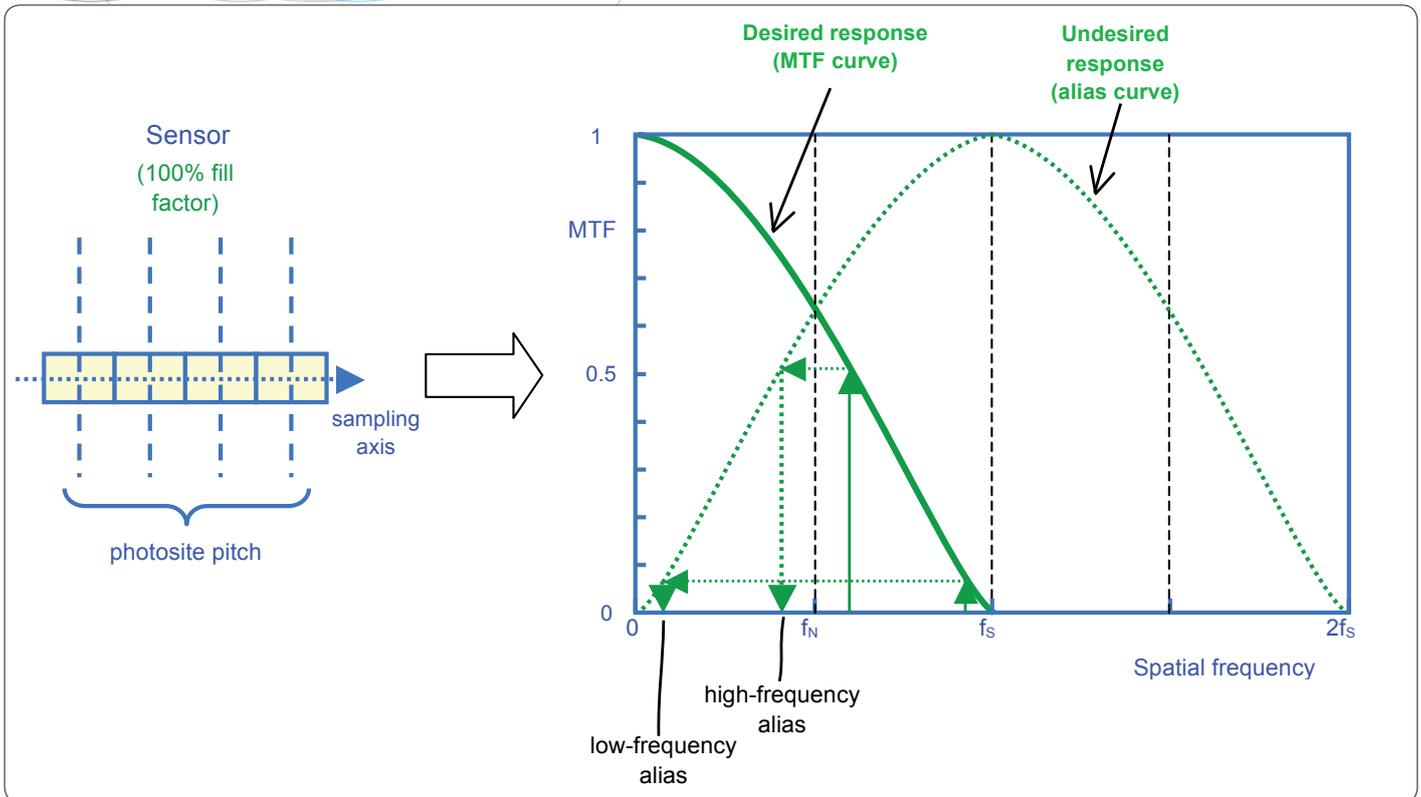


Figure 4 – Sensor with 100% Fill Factor

In Figure 4, the sensor now has 100% fill factor (touching photosites). The geometric MTF (solid curve) of this layout is now lower (63%) at f_N than with the 50% fill factor layout, but we can see too that the undesired alias amplitude is also lower at f_N and, most importantly, decays very rapidly towards zero amplitude at zero frequency. In effect, the greater window integration effect of the 100% fill factor is giving us a „free“ optical low pass filter, producing a valuable sharp cutoff of the most visible low-frequency aliases.

In practice, a fill factor of exactly 100% is not possible, since some degree of the total surface area of the sensor has to be taken up with non-light-sensing functions (e.g. transfer of electron charges from photosites to the output amplifier), but in a good design close to 100% is the aim and can be achieved.

Figure 5 shows that the MTF of the film in conjunction with the camera lens, the MTF of the scanner’s projection lens and the geometric MTF created by the layout of the pixels in the scanner’s sensors are all convolved (multiplied) together in determining the effective MTF between scene details and the digital data captured by the film scanner. Nyquist frequency in the scanner corresponds to approximately 80 lp/mm on the film.

So, is 4K the Optimum Scanning Resolution?

What the above sections initially suggested was that if a film scanner were designed to capture all resolution up to the ISO 12233 limiting resolution, alias-free, in the most extreme cases, such a scanner could be calculated to require a digital resolution much higher than 4K, perhaps as much as 11K digital resolution. An 11K

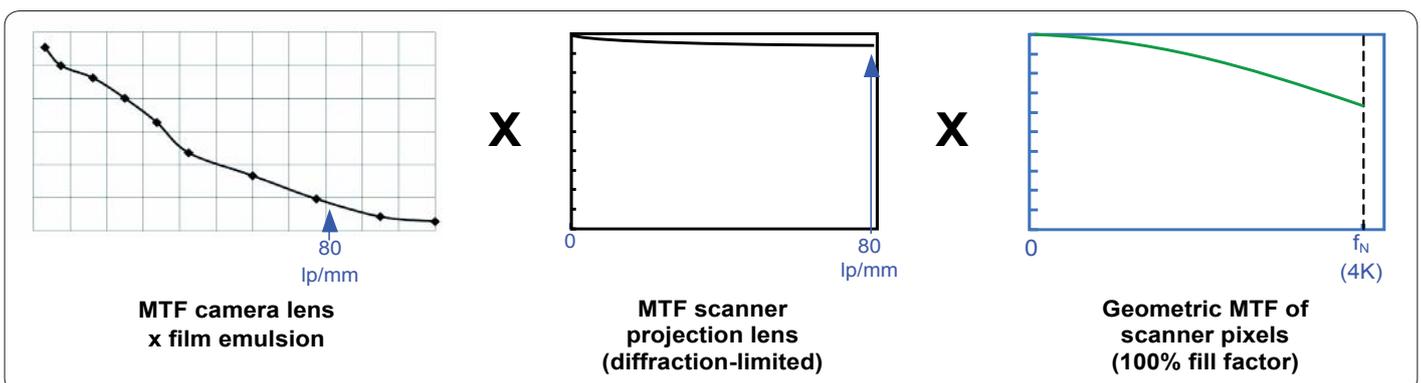


Figure 5 – Convolution of MTFs

scanner would be extremely expensive and slow in operation, and could have poor performance in other parameters, such as signal-to-noise ratio. In the majority of projects, the extra scanner pixels would not be capturing any additional image information compared to a lower resolution machine.

However, further examination of practical evidence indicates that provided the sensor layout is optimally chosen (close to 100% fill factor), all resolution up to the 80 lp/mm for OCN published by film stock manufacturers can be adequately captured with a 4K sensor architecture, because any aliases arising from unfiltered signal frequencies above this will either be:

- at very high frequencies and low amplitude and therefore not visible, or
- at lower frequencies and so reduced in amplitude by the pixel fill factor as to be invisible.

Scanning a first-generation OCN is also the extreme case. What the ITU tests also showed (Figure 2) was that after just one film generation (the answer print), the MTF fell to zero well before 106 lp/mm and even at 80 lp/mm was only about 4%; in fact, 20% modulation level was maintained only to about 50 lp/mm; this second generation's information content could therefore be captured adequately with far less scanning resolution than 4K; probably a 2.5K scanner would suffice.

Benefits of the 4K Design Decision

There are several benefits to limiting the information captured from 35mm film to that which is useful, i.e. scanning at 4K maximum.

The most important is the speed advantage. It is now possible to construct a 4K scanner that can run at up to 15 frames/second. Although this is achieved partly by the adoption of some special techniques – see below – such a speed would not be possible at a higher resolution of, say, 6K or 8K, because of the limited electron charge integration capability of the much smaller sensor photosite area (one quarter the area, assuming 8K versus 4K and equal fill factors). Conversely, attempting to run an 8K scanner at the same speed as at 4K would severely degrade the signal-to-noise ratio, detracting from the theoretical resolution benefit in the overall subjective assessment.

Another advantage of the 4K decision is the reduction in hardware costs. This arises not only from the sensors themselves, but also from the elimination of high-bitrate downconversion hardware, because few post-production workflows can handle 6K or 8K native scan data (those

that can handle it have to accommodate enormous data volumes - up to 300 MB per frame or more, compared to 75 MB for an equivalent frame at 4K*).

Interim Summary of the Scanning Requirements

1. Assume first generation 35 mm OCN film, exposed via a high quality production camera lens, and design for this challenging but practical case (leading to a 4K scanning resolution design decision).
2. Minimize scanner lens MTF loss with small-aperture low-aberration optics with large depth of field.
3. Employ close-to-100% fill factor pixel layout in the sensors for best separation ratio between wanted signal recovery and aliases.

SCANITY™ – A Realization of these Design Principles

A full description of the SCANITY™ 4K film scanner is given in other papers, but in this paper some of its features will now be briefly covered to explain the practical realization of the principles described above.

Film Illumination Source

The prior discussion considered scanning resolution in isolation. However, the relationship between resolution capability and illumination is very direct. Scanning resolutions have become higher with the result that for a given illumination level, the size of electron charge created in smaller and smaller pixel areas diminishes in inverse square proportion. Since quantum efficiency does not appreciably change, the only options then are to increase the effective charge integration time of the pixels (more on that later) or to increase the illumination intensity.

The density spectra of a colour film's emulsion layers require broadband illumination of sufficient intensity. To complicate matters, the camera negative has an orange mask (for color gamut enhancement purposes) that attenuates the incident light reaching the blue recording layer. At the other end of the spectrum, transmission of the red component of the illumination must be maintained to long wavelengths to read the modulation in the red recording layer, while not allowing harmful high-energy infrared to impose excessive heat on the film. Other film stocks, such as print, optimally require a different spectral distribution in the illumination.

Two techniques in the illumination method deal with these challenges. The first is that the illumination comes from the combination of discrete clusters of spectrally-

* Assuming frame sizes of 8192 x 6224 pixels (8K) versus 4096 x 3112 pixels (4K), and 16-bit data in R, G and B.

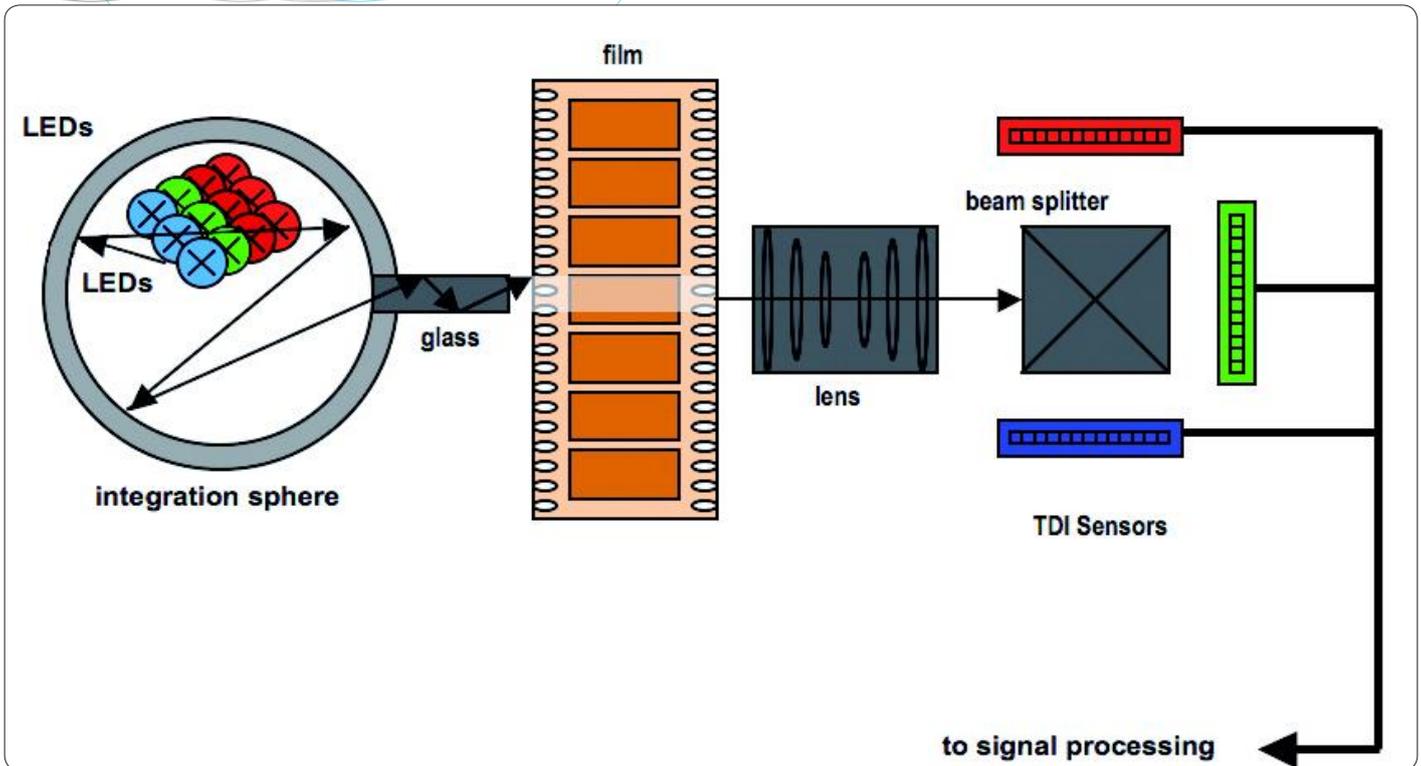


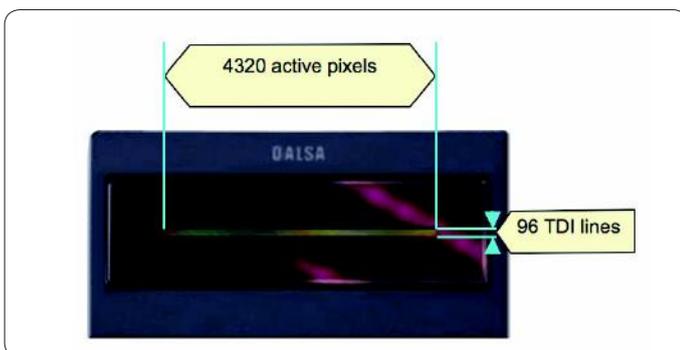
Figure 6 – SCANITY™ Light Path

separate red, green and blue LEDs (Figure 6). This allows accurate tailoring of the overall spectral power distribution to the density spectra of the emulsion layers, with a minimum of stray energy at unwanted wavelengths. A further refinement is that there are in fact two sets of red LEDs of slightly different center wavelengths; the appropriate set is used according to the type of film stock. After passing through an integration sphere, the light passes through the film via a very narrow slit and onto the sensors via a color beam splitter.

Sensor Architecture

The second technique extends the total integration period considerably, yet without slowing down the scanning speed (frames per second). This apparent contradiction is made possible by the concept of multiple timed charge integration periods in the scanner's image sensors.

Figure 7 – TDI Sensor



TDI (Time Delay Integration) architecture makes use of multiple repeated lines of pixels in SCANITY™'s linear array sensors (Figure 7); the same horizontal row of film „pixels“ is tracked as it moves through the gate by shifting the sensor's clocking from one sensor line to the next and turning on the LEDs each time for a brief burst of less than one scan line's time duration.

There are in fact 96 such lines in the TDI structure, so the same row of film „pixels“ is sensed 96 times as one TDI line after another is activated. At the same time, the integrated electron charges from prior lines are added to the current TDI line, causing an accumulation of charge volume. By the time the last TDI line has completed its sensing, the accumulation of charge is sufficient to multiply the voltage at the CCD output amplifier approximately 50 times. This gives an enormous improvement in effective sensitivity, and therefore signal-to-noise ratio, but without slowing down the transport speed or applying excess heat energy to the film. SCANITY™ makes TDI work successfully by precise synchronization between sensor line clocking, LED cycling, and the longitudinal positioning of the film in the gate by the transport servo.

Accommodation to Different Film Formats

Compared to area arrays, SCANITY™'s linear array TDI sensors have the advantage of adjusting automatically to the different frame heights of various film image formats, maintaining the same resolution for all of them. This principle works independently of the multiple TDI line structure in the sensors.

