

EBU - Tech 3335 : **Methods of measuring the imaging performance of television cameras for the purposes of characterisation and setting**

Alan Roberts, November 2013

SUPPLEMENT 010rev1: Assessment of a Canon EOS-1DC DSLR

Tests have been conducted in line with EBU R.118. This document is a report of the results of the tests defined in Tech3335 and is not an endorsement of the product.

The Canon 1Dc is a large, full-size, digital single lens reflex camera. As is becoming common, it also offers HDTV recording. Although there are few of the usual controls that go with video cameras, it was thought worthwhile to check its performance as an HDTV camera. The test procedures were as described and recommended by the EBU, in Tech 3335 (<http://tech.ebu.ch/docs/tech/tech3335.pdf>). The results were encouraging.

Initial tests were done on a camera body (DS126451, serial number 023011000227, a pre-production model) with two lenses (Canon 27-70mm 1:2.8 L II USM and Canon EF 70-200mm 1:2.8 L IS II USM). Later tests were also done on a production model (serial number 023011000037).

The camera body is big (158x164x83mm) and heavy (1.35kg), for a DSLR. Power is from a lithium-ion battery, rated at 11.1V/2.45Ah. A mains adaptor is available to replace the battery for extensive shooting. Recording is on to 2 Compact Flash card slots (high-speed, UDMA7 is required for 4k recording).

Recording video is QuickTime format (MOV) at 4k (4096x2160 at 24p, and 1920x1080, 1280x720, or SD , all at a variety of frame rates including 25Hz). Recorded files are in QuickTime (MOV) using Motion JPEG for 4k or H.264 for HD and SD (AVC, with long- or short-GoP, i.e. I-frame only). There is no facility for external time-code or genlock. Real-time video output is via HDMI, which appears to be of the same quality as video recording when recording HD or SD. Therefore it is theoretically possible to make better quality recordings externally than can be done via the in-built system. There is a small built-in microphone (adequate for note-taking and guide-sound recording) and an external mic socket (3.5mm). Connection to a computer via Ethernet allows for rapid file transfer, although it is probably quicker to swap cards for downloading. Hot-swapping caused no problems.

External controls are by push-buttons and rotary switches. Most of the controls are to do with stills-shooting, and have no effect on video-shooting. The menu structure, although quite comprehensive, has few items specific to video shooting. Therefore, in this document I have not listed the full menu contents.

This revision corrects one typographical error.

EBU - Tech 3335 : **Methods of measuring the imaging performance of television cameras for the purposes of characterisation and setting**

Alan Roberts, November 2013

SUPPLEMENT 010rev1: Assessment of a Canon EOS-1DC DSLR

The 1Dc is a full-size Digital Single Lens Reflex camera. The body under test offered recording facilities at 4k (at 24p) and HDTV at 1920x1080 (progressive at 50p, 25p and 25p), plus 1280x720 (progressive-only at 50Hz) and 640x424 (progressive-only at 25Hz). Monitoring and control of the camera are very different from a conventional video camera, making it rather unsuitable for many video purposes. The small LCD panel display was adequate for framing shots, and for getting focus provided the depth of field is short. Otherwise, it is better to use an external monitor fed via HDMI (Menu > Camera > HDMI output + LCD) when 'Live view' is selected (there's a button on the body for this).

The camera's single Bayer-patterned CMOS sensor is 35mm full-size, 36x24mm and has approximately 18.1 million photo-sites. The specification lists the choices of resolutions for stills-shooting, the maximum of which is 5,184x3,456, which is 17,915,904 pixels. Allowing for, say 22 blanked photo-sites along each edge (masked for use in establishing a black-level signal), the overall sensor size could be about 5,206x3,478, a total of 18,106,468 which is close enough to 18.1 million. Accordingly, the size of each photo-site must be about 6.9µm square, roughly double the area of the pixels of a conventional 3-sensor HDTV camera of nominal $\frac{2}{3}$ " size (5µm square). Thus, the sensitivity or noise should be about one stop better. The image-size changes when changing between stills and video, or between video standards, so the video feed must be taken by cropping the sensor for some modes. It seems likely that simple relationships have been chosen, using less than the full sensor, but in practice this should hardly matter since the image is nicely oversampled, but practical limitations often do not allow ideal performance.

For 4k movie shooting, there is a crop factor of about 1.3, implying that the image comes pixel-mapped from an area of the sensor exactly 4096x2160, and this implies that the maximum image-content resolution will be about $\frac{2}{3}$ of that, about 2880x1620. Note that this image format is not 16x9, it is actually 256x135, and that it should be cropped to 3840x2160 for television use. For Super-35mm shooting, the crop factor is stated to be about 1.6, and this implies that the image comes from an area of the sensor about 3200x1800 (which would then need a down-scaling by 5:3 to get to 1920x1080). For 'normal' 1920x1080 shooting, the image size appears to about full size, thus implying a down-scaling by about 2.7:1, 8:3.

All testing was done by recording onto Compact Flash in the camera, since this is the mode in which the camera is easiest to use and will probably be used. Live monitoring was done on a 42" plasma display via HDMI. The test procedures were as described and recommended by the EBU, in Tech 3335 (<http://tech.ebu.ch/docs/tech/tech3335.pdf>).

The prime purpose of this test was to establish the best shooting mode for use in HDTV. Inevitably, some of the more traditional aspects of camera testing received less than the usual attention.

1 Colour performance

The camera menus provide control over the system primaries (Menu button > Camera menu 1 > Color Space) offering sRGB and Adobe RGB. Since the sRGB primaries are those of ITU.709, the world-wide HDTV standard, only this setting makes sense for shooting video. Effectively, this setting is a choice of matrix rather than of primaries.

There are also Picture Controls (Menu > Camera menu 1 > Picture style) offering a choice of 7 preset settings and 3 user-definable. 'Standard' was used for all the tests. However, by then pressing the 'Info' button, the stored settings can be modified (Sharpness, Contrast, Saturation, Color tone). Most of these settings were left at the default, factory value, only Sharpness was explored. However, the gamma curve can be changed (Menu > Camera menu 5 > Canon Log settings) so this was investigated briefly.

Using sRGB and Standard Picture Control, the colour performance was perfectly acceptable on a television display. Colour-rendering was generally good, no individual colour standing out as being problematic. The



Figure 1 Colorchecker (a) sRGB gamma

(b) Canon Log

Canon-Log curve clearly needs post-processing to get the best from the pictures, as expected. However, the apparent ‘speed’ of the camera drops significantly in this mode, the white patch is at only 74% of the luma-signal level of the sRGB image, effectively about 1.5 stops. This is a reasonable indication that the log curve can cope with at least 2 stops more exposure, as headroom, quite apart from the obvious extra ‘foot-room’. When using the log curve, all capture and post-processing should be done to at least 10-bit precision in order to preserve fine colour gradations.

2 Video Resolution

The camera was exposed to a circular zone plate test chart, containing patterns to test luma, R G and B, and chroma channels. In each case, the chart was framed to fill the image width on the live-monitoring 1920x1080 feed.

2.1 Video Resolution at 4k

Fig. 2 shows one quadrant of the luma pattern of a circular zone-plate test card, captured at 4k.

The smaller pattern (bottom right) explores the 4k frequency range, which could be read as either 4096x2304 or 3840x2160, depending on exactly how the chart is framed to fill the live-monitoring feed of 1920x1080. If the chart is framed to fit the image height, it is 4096x2304. If it is framed to fit the width, it is 3840x2160 and thus the larger quadrant explores the frequency range of HDTV, 1920x1080.

The smaller pattern shows the usual diamond-shaped resolution limit, a function of the Bayer-pattern layout of photo-sites. There is coloured aliasing at high horizontal and vertical




Figure 2 Resolution, 4k capture

Tests have been conducted in line with EBU R.118. This document is a report of the results of the tests defined in Tech3335 and is not an endorsement of the product.

frequencies, and luma aliasing at high diagonal frequencies. However, all these aliases are at a fairly low level, indicating that there is a good optical spatial low-pass filter, matched to the 4k mode.

The HDTV quadrant shows no aliasing, and clean resolution up to the pattern edges. This is the resolution which the camera must attempt to deliver in the HDTV capture modes.

2.1.1 Video Resolution at 4k, down-scaled to 1920x1080

For this test, the captured video clip was ingested into Edius 6.5, and the project set to 1920x1080/24p. Although 24p is not a frame rate for television, this was chosen to avoid any temporal conversions to 25p.

A still frame was exported using the Lanczos 3 algorithm, Fig. 3a. The HD resolution is well-maintained with very little aliasing, and most of the extra of the original 4k image is well suppressed, the alias levels in the 4k quadrant are quite low. This is quite a good down-conversion.




Figure 3 Resolution, 4 k capture (a) Lanczos 3 scaling

(b) conventional scaling

The Lanczos 3 algorithm is believed to be the best available for this purpose, and is available in all high-end video-editing software. If it is available, it should be used for preference. It is a much more complex filter than the alternatives and requires significantly more computing power to accomplish, but the results are well worth it.

For comparison, Fig. 3b shows the performance of the alternative, simpler, scaling available in Edius 6. The HD resolution is not quite so well maintained, and the high-frequency aliasing in the 4k quadrant is significantly less suppressed.

2.2 Video Resolution at 1920x1080

There are several capture modes at 1920x1080, falling into two groups called S35 and 1920. The performance is quite different in these two modes.

2.2.1 Video Resolution at 1920x1080, S35 capture

Fig. 4a shows the luma quadrant with S35 capture, using the 'All I' mode, which is M.264 (MPEG4) using only I-frame compression. There is a significant loss of resolution at high frequencies, comparing with the performance when captured at 4k (Fig. 3). This is hardly surprising since the image content comes from an area about 3200x1800 rather than from the full frame, and thus the lens focal length appears to have been increased by about 1.6:1.




Figure 4 Resolution, S35 capture (a) All-I

(b) IPB

The ‘extinction’ frequencies are about 80% horizontally (1530) and vertically (864), which may not be really good enough for HDTV at the high-end of production. High-frequency content in the 4k, smaller, pattern is well suppressed, indicating that the down-scaling is being done with proper filter-based conversion rather than simpler interpolation.

Fig. 4b shows S35 capture using the ‘IPB’ mode, i.e. Long-GoP M.264. The performance is almost identical, but not quite, the level of aliasing is a little higher, and the clean resolutions limits a little lower. It is probable that the same filtering has been used, but that the frame grab has been made from an interpolated frame rather than an I-frame. Nevertheless, the performance can be regarded as almost identical.

2.2.2 Video Resolution at 1920x1080

Fig. 5a shows the luma quadrant with ‘1920’ capture, using the ‘All I’ mode, which is again M.264 (MPEG4) using only I-frame compression. There is now quite a dramatic loss of resolution at high frequencies, comparing with the performance when captured in S35 mode 4k (Fig. 4). This is quite surprising since the image content comes from the full frame.




Figure 5 Resolution, 1920 capture (a) All-I

(b) IPB

The ‘extinction’ frequencies are about 66% horizontally (1300) and vertically (730), which is certainly not be good enough for HDTV at the high-end of production. High-frequency content in the 4k, smaller, pattern shows some worrying aliasing, indicating that although the down-scaling may be using proper filter-based conversion rather than interpolation, the filter isn’t sufficiently large or complex.

Fig. 5b shows 1920 capture using the ‘IPB’ mode, i.e. Long-GoP M.264. The performance is different, a little sharper horizontally and less sharp vertically. Also, the aliasing in the 4k pattern is different. I have no explanation for this, except that the grab has probably been made from an interpolated frame rather than an I-frame. Nevertheless, the performance of both modes is inadequate for high-end HDTV production.

2.3 Video resolution, sharpening

Unlike in conventional cameras, the 1Dc has only one control for setting ‘sharpness’. This there is unlikely to be a setting best suited to television use (since it has probably been developed for stills shooting).






Figure 7 Sharpening, (a) sharpening off (=0)

Nevertheless, it works in video capture modes and so will inevitably be used.

(b) maximum (=7)

Fig. 6a shows the luma quadrant (in 4k mode, captured as 1920x1080 using the conventional scaling algorithm of Edius 6) when sharpness was set to zero. This is the native resolution of the camera, and it is significantly softer than in Fig. 3b which used the default sharpening setting.

Fig. 6b shows the effect of the maximum sharpening setting. Although the detail level has increased significantly, the major effect is to enhance contrast at lower frequencies. Also, there are new alias null centres at about 84%, caused by harmonic distortion of the sinusoidal modulation of the test chart, and there is a significant increase in the alias levels in the 4k pattern.



Since the only sharpening control is the level, it is not

Figure 6 Sharpening, level=2

possible to tailor sharpening for high-end HDTV shooting. The default level (3) is too high, level 2 appears to be a little better, Fig. 7. However, it is probably preferable for all sharpening to be turned off for HDTV shooting (if 4k shooting is used and Lanczos 3 used for down-scaling), and any sharpening thought to be necessary to be applied in post-production where more sophisticated controls should be available.

3 Sensitivity

To measure the sensitivity, the camera speed was set to ISO 400 with the shutter set to 1/50 second, as will probably be used for HDTV shooting. It was exposed to a pair of Kodak Gray cards, to show both the white (90% reflectance) and mid-grey (18% reflectance) simultaneously. With a conventional video gamma-correction curve, the grey card should produce about 50% video level when the white side produces 100% video. In this case however, the white card produced rather less than 100%, indicating that the gamma-correction is not standard. Therefore, rather than measure the light level at which the white card made 100% video, I measured the level at which the grey card produced 50% level, which is more indicative of typical exposures.

Illuminating the two cards at 2000 lux, the grey card made 50% video signal at F/8. So the sensitivity is nominally F/8 at 2000 lux.

4 Noise

Many measurements of noise were made on the production model, in the light of comments received about the camera’s performance, in order to assess the effects of 4k and S35 shooting.

Fig. 8a shows the noise profile at 800 ISO, using 4k capture. Noise-reduction was left in the ‘Std’ setting. This mode records frames individually, intra-frame recording. The solid lines with diamond markers are measurements of the 4k capture, those with dotted lines and cross markers are measurements of the same captured file down-scaled to 1920x1080 using the Lanczos 3 algorithm. Clearly there is a noise reduction of between 1 and 3dB resulting from the down-scaling.




Figure 8 Noise at 4k (a) profile at ISO 800

(b) noise vs ISO

The curves are normally expected to continue upwards as signal level decreases, since noise generated in the sensor is differentially amplified by gamma-correction. Clearly this effect is severely limited below about 40% luma level. This seems to imply that there is differential noise-reduction. It is worth noting that this odd profile in 4k mode makes the pictures look less noisy than the measurements would imply. This is because visual assessment of noise is generally dominated by the noise near black, which in the 4k mode at least, is highly suppressed. The general appearance is approximately 6dB better than the numbers indicate.

Fig 8b shows the distribution of noise at mid-grey, also in 4k capture mode. The brown curve is measurements of the 4k-captured files, the cyan curve is measurements of the same files down-scaled to 1920x1080 using Lanczos3, again showing the same noise advantage. Clearly, shooting at 4k and down-scaling to 1920x1080 in post-production using Lanczos 3 is a good plan.

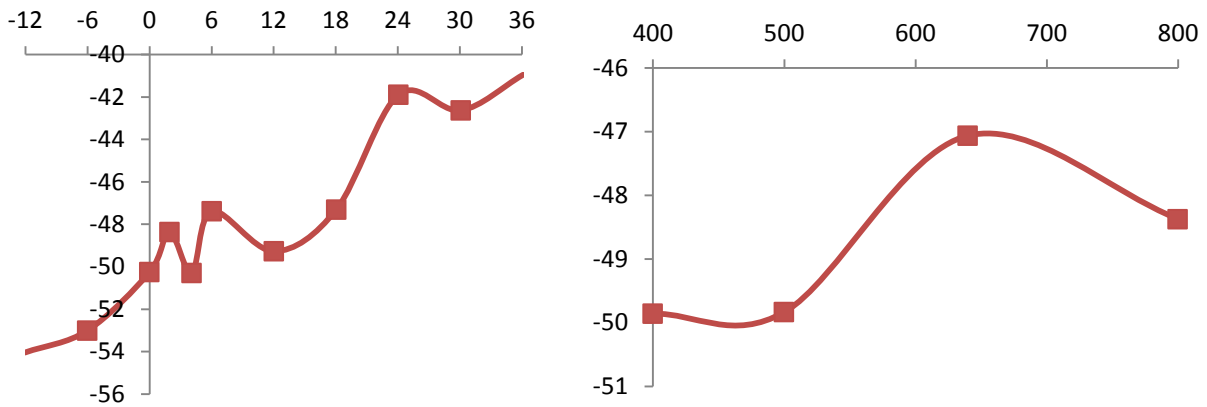


Figure 9 Mid-grey n vs 'gain' (a) S35-IP, NR=std (b) S35 ALL-I, NR=off

Fig. 9a shows the distribution of noise at mid-grey, over the full range of camera 'speed' settings. In this case, the horizontal axis is marked in video terms, dB, relating to the default value of 400 ISO as the zero-gain point. The general trend of this curve should be a straight line with a slope of 1/2, i.e. a worsening of the noise level by 3dB for each increase in gain of 6dB. The S35 line does indeed follow this trend (except for one rogue value at +4dB or 640 ISO) which is may not be significant.

Fig.9b is a close-up of the noise levels in S35 ALL-I capture mode, with noise reduction switched off. It establishes that the general trend for increasing noise level with gain does apply, and that the odd value at 640 in Fig.9a might be an anomaly. Alternatively, it could be evidence that there is something odd in the signal processing, indicating that noise levels are lower than expected at multiples and sub-multiples of 640 (i.e. 160, 320, 640, 1280 and so on). It would require much more exhaustive tests can establish this for certain.

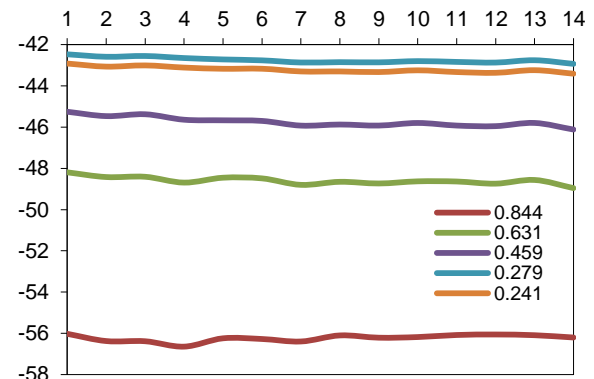


Figure 10 Noise level in 14 successive frames, S35-IP capture

Fig.10 shows the noise level in S35-IP capture mode. This is standard MPEG2, and so there is a possibility of the noise level being affected by the GoP structure. The curves are for five levels of the luma signal, spanning 14 frames so that there must be two I-frames.

Clearly there is no significant change of noise level which could be correlated with the GoP structure.

4.1 Noise reduction

All the previous measurements were made with the Noise Reduction item in the menus disabled or in Std setting. Fig. 11 shows the effect of noise reduction. In the menu, there is a comment that noise reduction is most effective at the highest camera speed settings, so tests were done at 12800 ISO. There is an improvement of about 2dB across the scale, and the reduction does not appear to do significant damage to resolution.

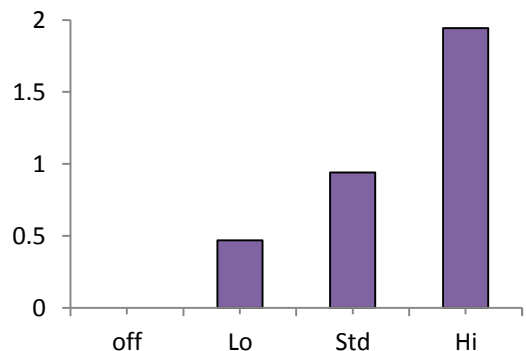


Figure 11 Noise reduction at ISO 12800

Since the effect of noise reduction at low ISO settings (low gain) is minimal, there is no penalty involved in using noise reduction. The Std setting performs well and there appears to be no reason for not using the Hi setting.

5 Exposure range

This is normally calculated as the ratio of the exposure which just causes white clipping to the exposure level below which no details can be seen. The short-form manual makes no mention of the exposure range (the main manual was not available during the tests).

Now, the Colorchecker chart has a contrast range of 28.76:1 (white to black) which is about 4.85 stops. The signal levels for the grey-scale in the captured image, using Canon Log gamma (Fig. 1b), run from 0.2306 to 0.6522. If the log curve follows the same curve up to peak white, then there is capacity to capture approximately a further 3.5 stops, and if it extends downwards, about another 4.5 stops, making a total of about 13 stops. This analysis is by no means rigorous, and will be limited by the noise level.

6 Motion portrayal, rolling shutter

Since this is a CMOS camera, the image data is probably read from the sensor by scanning, rather than by taking a very brief, global, reading of the values into a store for later scanning. This process, called 'rolling shutter', is the same as scanning in a vacuum tube camera or CRT television set, and can cause severe geometrical distortions when there is significant motion in the image.

To test for this, a small fan was set up, rotating at a speed designed to cause strobing of the fan blades. Fig.12 shows the results.

With the shutter set to 1/320, each frame shows a fairly sharp image of the rotating blades, and it is clear that the distortion due to the scanning process is quite small. The blade on the right, travelling downwards, is widened by the scanning process, that on the left moves against the scanning and is narrowed. At the top and bottom, when motion is across the direction of scanning, the edges are bent a little to the left. This effect is independent of the shutter duration.




Figure 12 Rotating fan

7 Infra-red response (IR)

No video camera should respond to IR. The simple test for this is to take a remote control unit and point it into the lens, then start pressing buttons. If the IR LED is seen lit or flashing, then the camera must be responding to IR since all consumer remote controls use LEDs which emit 'light' at about 930nm, well beyond the red extreme of the visible range.

This camera showed no response to IR. Any camera which does show a response to IR will show odd colour behaviour under some lighting conditions, and will potentially have unstable black level, due to this light pollution.

8 Conclusion

The camera has very limited controls when in video mode, but has reasonable connectivity, allowing full-resolution external monitoring and recording. Sound facilities are sparse, the internal microphone is adequate for note-taking or guide sound, and the microphone connector is a 3.5mm jack offering only un-balanced input.

The test procedures were as described and recommended by the EBU, in Tech 3335 (<http://tech.ebu.ch/docs/tech/tech3335.pdf>). Video performance is quite good at 1080p, and was not tested at 720p.

When recording in 4k mode, and down-converting to HDTV in post-processing, the resolution is exceptional, and the level of aliasing is very low. Best performance is achieved by using the Lanczos 3 algorithm. When recording in S35 mode, the pictures are rather soft, only reaching about 1530x864, but the level of aliasing is very low. When recording in 1920 mode, the pictures are decidedly soft, only reaching about 1300x730, and the level of aliasing is higher though not especially so. Detail enhancement (sharpening) is crude by video standards and would be better done in post-production.

There is no quality advantage to making external recordings via the HDMI output.

Noise levels are not quite as good as can be got in a conventional video camera, but are acceptable, i.e. it achieves the EBU R.118 high-end production limit of -50dB up to about ISO 500 in 4k capture mode, and the down-conversion to 1920x1080 gives another 2dB or so, provided the Lanczos3 algorithm is used. Noise reduction can further improve noise by up to 2dB, which would make ISO 640 or 800 acceptable. However, the noise profile is not as expected, indicating that some form of noise reduction is permanently applied, but this has the effect of making the pictures look a little less noisy than they really are, by perhaps 4 to 6 dB, implying that even ISO 1600 could be usable. Noise reduction works well, in that it works only when the noise levels increase as with increased ISO setting.

The camera does not do well enough in either of the S35 capture modes.

Colour performance is good, and the camera does not respond to infra-red illumination.

Motion portrayal is good, the effects of the rolling shutter are nicely suppressed.