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

Alan Roberts, May 18 2012

SUPPLEMENT 002: Assessment of a Nikon D4 DSLR

The Nikon D4 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 (serial number 2020941) with two lenses (Nikkor AF-S 20-70 1:2.4G, serial number 546609, and AF-S 24-120/4G VR ED, serial number 62001283).

The camera is big (160x156x90mm) and heavy (1.34kg), for a DSLR. Power is from a lithium-ion battery, rated at 10.8V/2Ah. A mains adaptor is available to replace the battery for extensive shooting. Recording is on to 2 card slots, one Compact Flash, the other XQD.

Recording video is at 1920x1080, 1280x720, or 640x424, at a variety of frame rates (including 25Hz). Recorded files are in QuickTime (MOV) with H.264 compression (AVC, MPEG4). 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. Therefore it is 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 or USB allows rapid file transfer, although it is probably quicker to swap cards for downloading. Hot-swapping caused no problems. The camera also has an optional wireless connection system.

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. Some controls can be assigned to preferred functions. 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.

Colorimetric and Resolution requirements of cameras

Alan Roberts (May 14-20 2012)

ADDENDUM 71 : Assessment of a Nikon D4 DSLR

The D4 is a full-size Digital Single Lens Reflex camera, with HDTV recording facilities at 1920x1080 (progressive-only at 29.97, 25, 23.976Hz), 1280x720 (progressive-only at 59.94, 50, 29.97 and 25Hz) and 640x424 (progressive-only at 29.97 and 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 filed is short. Otherwise, it is best to use an external monitor fed via HDMI (Setup tools>HDMI).

The camera's single CMOS sensor is 35mm full-size, 36x23.9mm and has approximately 16.8 million photo-sites of which 16.2 million are 'effective', i.e. used to make pictures (the remainder being blanked off to provide thermal and black level data). The specification lists the choices of resolutions for stills-shooting, the maximum of which is 4,928x3,280, which is 16,163,840 and close enough to 16.2 million. Allowing for, say 76 blanked photo-sites along each edge, the overall size could be about 5,004x3,356, a total of 16,793,424 which is close enough to 16.8 million. Since the image size did not change when changing from stills to video or between video standards, the video feed must be generated from most of the sensor, probably 4,928x2,772 using a 2.57:1 (77:30, an awkward relationship) down-scaling filter for 1080-line and 3.85:1 (77:20, similarly awkward) for 720-line. It seems likely that simpler relationships have been chosen, using slightly less than the full sensor, but in practice this hardly matters since the image is nicely oversampled. which should result in fairly simple conversions.

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. 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).

1 Colour performance

The camera menus provide control over the system primaries (Menu button > Shooting Menu > 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 (press the button with a key symbol) offering a choice of 6 settings. 'Standard' was used for all the tests. However, by then pressing the 'OK' button, the stored settings can be modified (Sharpening, Contrast, Brightness, Saturation, Hue). Most of these settings were left at the default, factory value, only Sharpening was explored.

Using sRGB and Standard Picture Control, the colour performance was perfectly acceptable on a television display.

2 Resolution (1920x1080)

The camera was exposed to a circular zone plate test chart, containing patterns to test luma, R G and B, and chroma channels. Only one quadrant of the luma pattern is shown here (Fig.1).

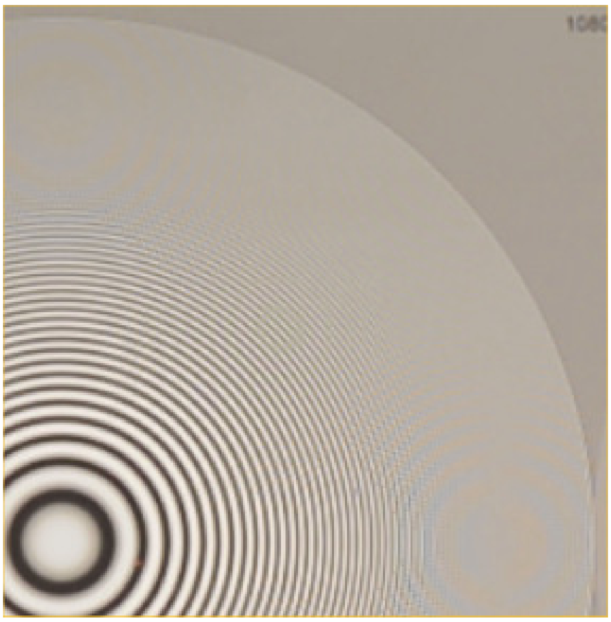

Figure 1 Resolution 1080p, default detail

This showed some low-level coloured aliasing, apparently centred on 1,643 horizontally and 924 vertically, indicating a 'three-ness' in the down-sampling (since 1,650=4,928/3). This is to be expected if the size of the down-scaling filter is not very big. However, the level of aliases is fairly low and might not be visible on most video sequences, but the resolution does not substantially exceed 1,355x762, and the presence of coloured aliasing implies that the resolutions of R and B are not the same as G.

The coloured aliases are the result of unequal RGB resolution in the Bayer-patterned sensor. The resolution chart clearly shows this (Fig.2). Note that the red pattern (and therefore also the blue pattern) has aliases centred on 1,643:0 (i.e 1,643 horizontally and 0 vertically) and 0:924 (0 horizontally and 924 vertically), but otherwise shows square resolution up to about 1,355x762, while the green pattern has aliases centred on diagonal frequencies 1,643:924. This is exactly as is expected for a Bayer-patterned sensor.

The resolution chart also contains two smaller patterns, which explore the frequency range up to 3,840x2,160, again, only one quadrant is shown (Fig.2 centre top). Again, there are aliases visible centred on about 3,180 horizontally and 1,790 vertically, confirming a 3:1 relationship with the sensor dimensions. Again, this is evidence that the down-scaling filter is not big enough, and is not rejecting the high-frequencies which the high-resolution lenses can deliver to the high-resolution sensor.

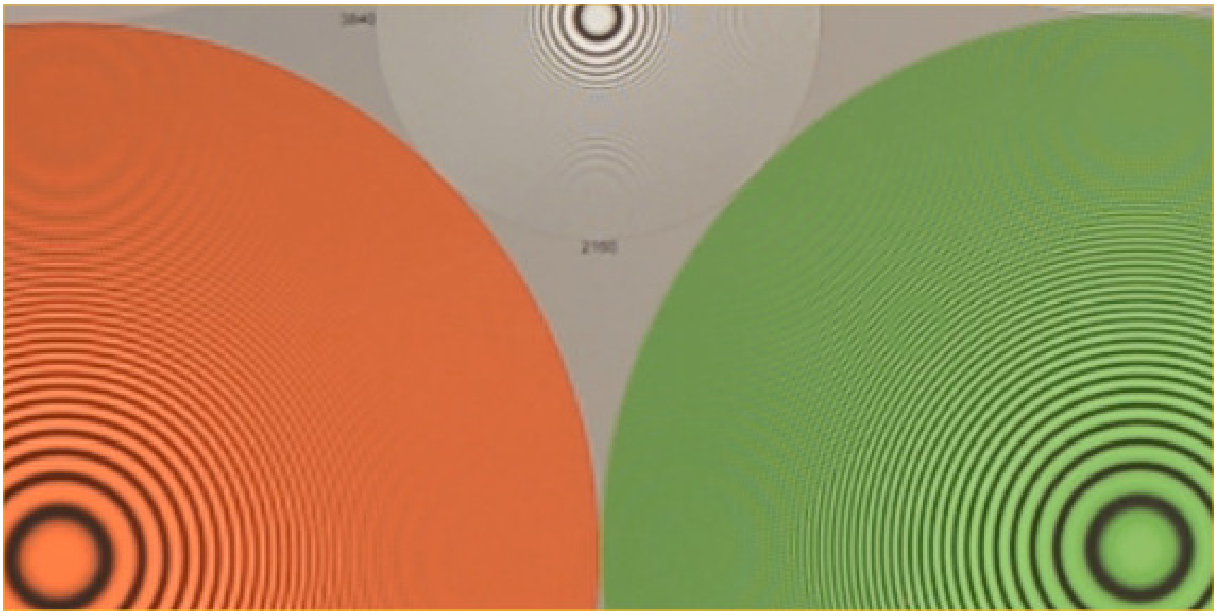

Figure 2 Resolution 1080p, default detail

Although the delivered resolution is not full 1920x1080, and aliases are visible, the performance is quite good for a stills camera. However, a proper video camera of the same price would normally be expected to perform rather better.

Spatial aliasing is caused by less-than-ideal interpolation, and in still images may not be a big problem, because they don't move. However, aliases in moving pictures are much more of a problem because, when the image moves, the aliased frequency content moves in the opposite direction to the image motion, causing a rippling effect on edges. Since motion-sensitive compressors such as MPEG2 and MPEG4 depend on the cleanliness of edges to measure motion, these aliases can cause the compressor to allot undue bit-rate to motion and/or result in excessive compression artefacts. Despite all this, pictures with aliasing at the levels seen here are probably acceptable in HDTV.

3 Resolution 1280x720

This time, the story is not so satisfactory (Fig.3).

The aliases centred on 1,643 and 924 are still visible, resulting from down-scaling to 1920x1080, but there are second-order aliases centred on about 740 and 415. These result from the further down-scaling to 720p, and could be more problematic.

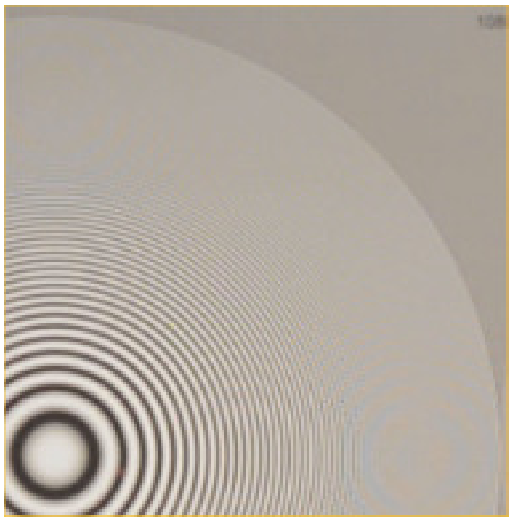

Figure 3 Resolution 720p, default detail

The resolution is clean only up to about 1,030x580. Given the low achieved resolution and the level of multiple aliases, this mode is not recommended for HDTV shooting.

4 Gamma curve

Normally, the distribution of noise levels matches the slope of the gamma curve, so it would have been good to be able to use a standard curve, but there were no such options in this camera. Therefore, several measurements were made of the patches of the grey scale on a Colorchecker chart (which has known reflectances) at exposures from F/2.8 to F/22, all at ISO 1,000 with shutter set to 1/50 second. Some slight manual adjustments in the exposure calculations were needed to get the points toile on the same curves nicely. Fig.4 shows the result, plotted on linear and logarithmic axes.

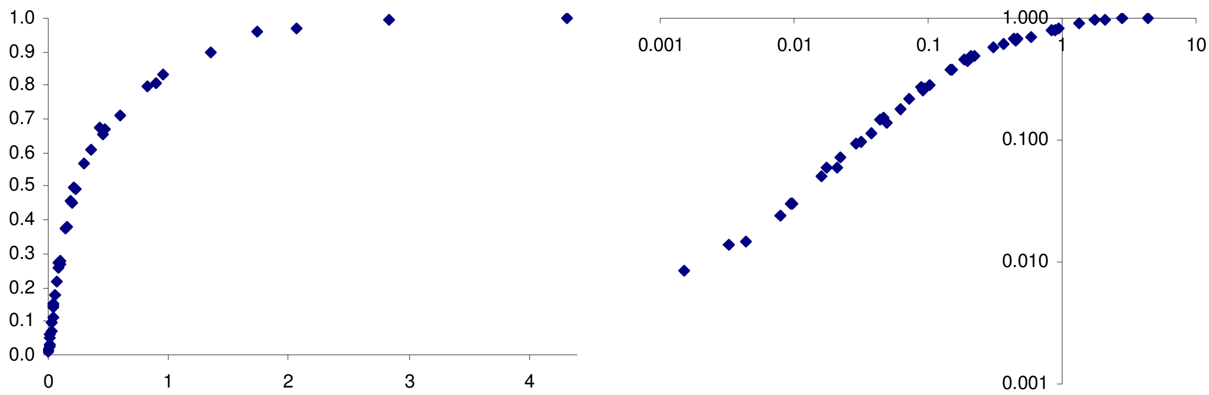

Figure 4 Gamma curve (a) linear axes **(b) logarithmic axes**

The 'normal' exposure was chosen such that 100% exposure caused about 85% signal level, although this is purely arbitrary. In the log plot, it is clear that the logarithmic nature of the curve expires at exposure level 0.006, which is very low for a video gamma curve. Also there is considerable compression of highlights above about 0.2 exposure, but the changes of curvature are gentle at both extremes. This is typical of proprietary 'log' curves used in stills cameras, and in cameras intended to give a filmic look to video shooting.

5 Noise measurements

As a by-product of the gamma curve measurement, it was possible (although tedious) to get many measurements of noise level, at many different signal levels, Fig.5 shows the results.

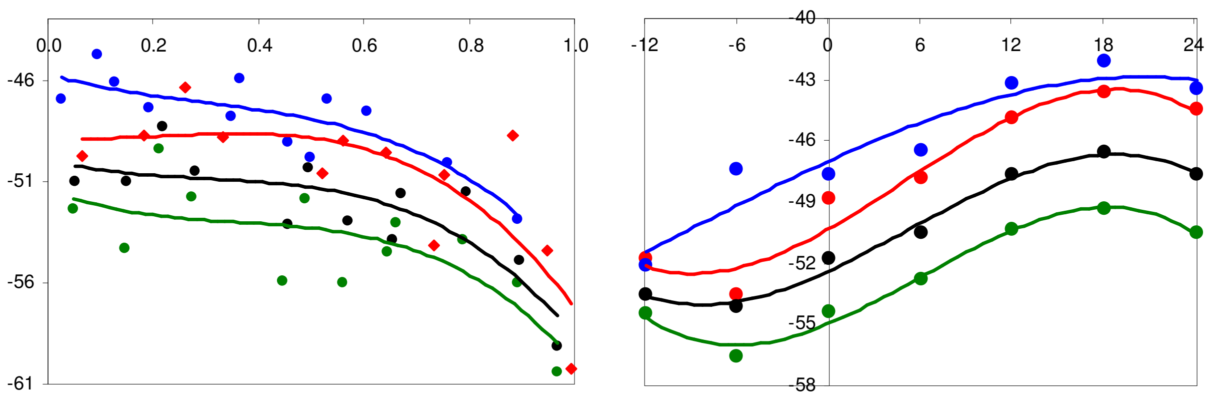

Figure 5 Noise 1080p (a) at ISO 1,000 **(b) at about 50% video level**

With the camera set to ISO 1,000 and 1/50 shutter (Fig.5a), the distribution of noise versus signal levels is interesting but not unusual. The uncertainty in each noise measurement is about ±1dB because only a small part of the video signal was used (about 60,000 pixels) therefore the points are not joined into lines, instead, trend lines are plotted. Nevertheless, the noise levels are consistently low, around -54dB for the luma signal, which is widely accepted as a 'good' level for HDTV.

The curve would normally be expected to follow the slope of the gamma curve. The decrease in noise near white is probably due to the high compression of dynamic range near white, and the limited rise in noise near black to a limitation in the gain-bandwidth product of the analogue circuits on the sensor.

Varying the 'speed' from ISO 200 to ISO 12,800, and plotting measurements for signals at about 50% video level (Fig.5b) shows the expected change of noise levels of about 3dB/stop. The horizontal axis is marked in terms of video gain (dB), where 0 is set for ISO 800 (so ISO 200 is -12dB, ISO 12,800 is +24dB). Again, there is an uncertainty of about ±1dB due to the use of relatively small areas of the image for measurement. The indications are clear, even at +18dB gain or ISO 6,400, the noise levels are still acceptably low. Even at ISO 12,800 noise may not be a problem. The reduction of noise level at ISO 12,800 is probably due to limitations in the gain-bandwidth product of the amplifiers providing the gain. The points are not joined into lines because they come from many disparate measurements; instead, trend lines are plotted. The noise levels are remarkably low.

Since the photo-site spacing is approximately 36/4,928=7.3µm, the maximum area for each photo-site is about 53.4 μm^2, the camera should be about 1 stop 'faster' than a 3-sensor ⅔" camera, which is normally expected to deliver better than 50dB PSNR at it's zero-gain setting. Such cameras are normally rated at between ISO 640 and 800, and so we should expect either that this camera is 6dB better for noise, or delivers the same noise at a stop faster. The results appear to confirm this supposition, approximately.

6 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. From the gamma-curve plot, the over-exposure range is about 4.5 times, 2.2 stops, and the gamma curve appears not to break off into a linear portion near black, it more closely resemble a film-type logarithmic curve. Therefore, the only limit on low exposure is the noise level. To confirm this, I took multiple exposures at ISO 1,000, and looked at the noise levels.

Exposures from F/2.8 to F/22 were all almost acceptable. At F/2.8, the top two patches of the Colorchecker chart were both clipped, but at F/4, all patches were clear of clipping. Therefore the clipping level appears to be at about F/3.5. At F/22, the noise level near black (3% signal level) is about -51dB, a remarkably low level. Raising this signal level by 2 stops in a video editor, produces 12% signal level, clearly visible with detail, and noise at about -39dB, still acceptable.

Now, the Colorchecker chart has a contrast range of 28.76:1 (white to black) which is about 4.85 stops. Add to this the 6 stops of lens exposure used in the tests, and the 2 stops which can be comfortably stretched in post-production, the total exposure range comes to about 13 stops, which is truly remarkable.

7 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, 6 black sectors on a light disc. Fig.6 shows the results.

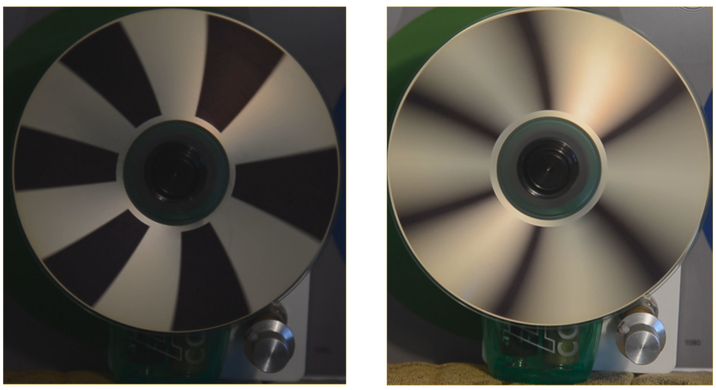

Figure 6 Rolling shutter (a) 1/1,000 (b) 1/100

With the shutter set to 1/1,000, each frame shows a sharp image of the rotating blades, and it is clear that the distortion due to the scanning process is quite small. Blades on the right, travelling downwards, are widened by the scanning process, blades on the left move against the scanning and are narrowed. At the top and bottom, when motion is across the direction of scanning, the edges are bent to the left. This effect is independent of the shutter duration.

8 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 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.

9 Camera settings

The camera menus have few items which directly affect the video imaging performance. Only Sharpening was investigated, since all the other controls which modify the image appear to be special effects rather than

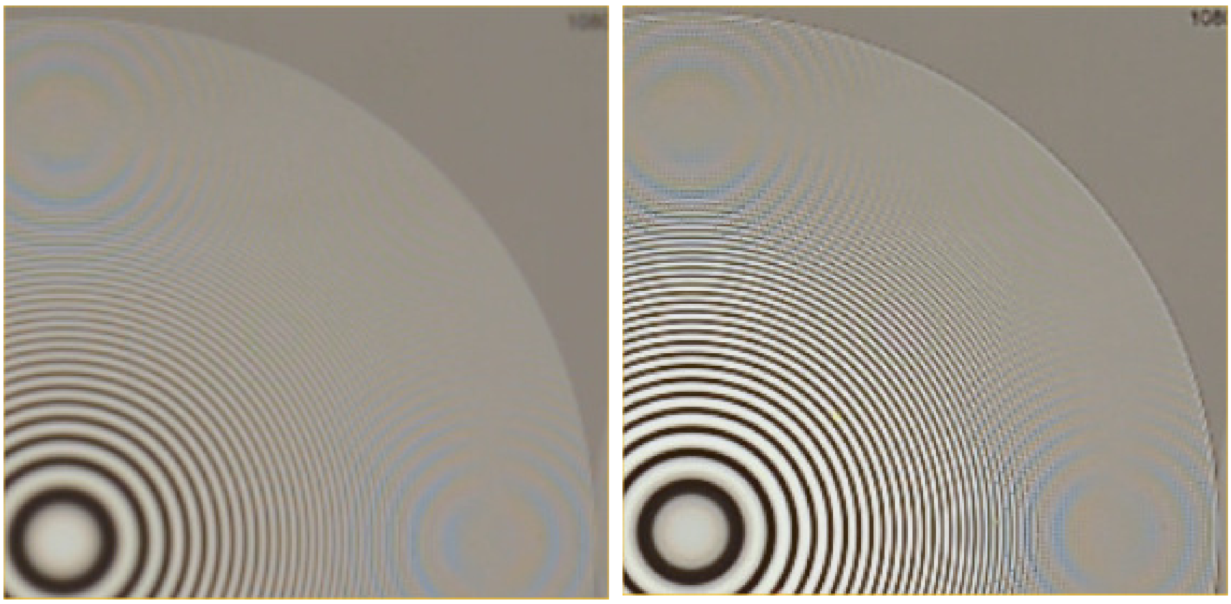

Figure 8 Sharpening (a) 0 (minimum) (b) 9 (maximum)

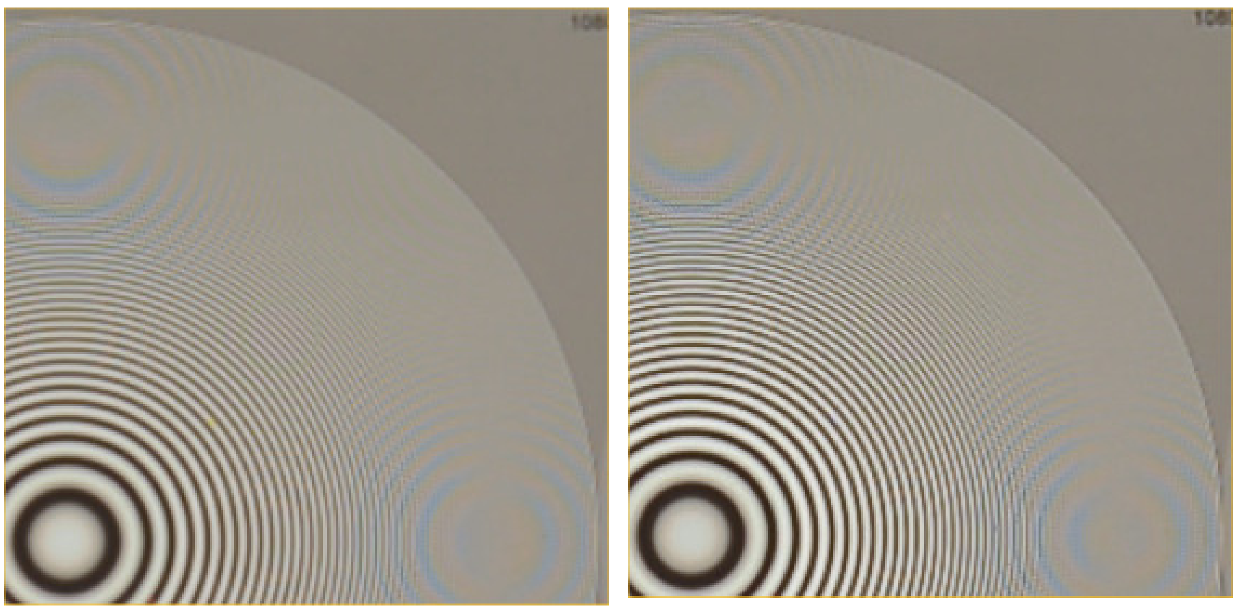

Figure 8 Sharpening (a) level 2 (b) level 4

the normal image controls found in conventional video cameras.

With Sharpening set to 0 (minimum value, Fig.7a), the image is decidedly soft. Setting Sharpening to 9 (maximum value, Fig.7b) greatly enhances contrast at low frequencies and emphasises the coloured spatial aliasing, neither of which is desirable. Since there are no controls other than the level, nothing can be done other than to select a setting which is acceptable.

The lower values are more promising. Fig.8 shows levels 2 and 4. Of these, the lower obviously emphasises the aliasing less and is more 'restful'. Level 4 should be regarded as the absolute maximum value for video shooting, with level 3 or 2 used for preference.

10 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, much less so at 720p. However, even though the sensor has 16.8 million photo-sites, it achieves only about 1,355x764, only a little better than 1,280x720.

Noise levels are very low, and the camera could be acceptable with ISO settings up to 6,400. Coloured spatial aliasing is present, and is clearly visible even on the camera's LCD display (921k pixels, about 1,176x784). In theory, meticulous control of the shooting style can minimise this, by using only motivated pans together with fairly short depth of field, such that detail out of the focused plane is always soft and therefore can never provoke aliasing. Exposure range is, potentially, as high as 13 stops, although this will be limited by the acceptability of the noise levels near black.

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.