Title: CMOS Waveform Distortion (AKA: “DVXuser Mythbusters: The EX1 looses resolution when panned.”)

Synopsis.
Waveform monitors are the main tool for measuring video. But when there are rolling shutter distortions from a CMOS sensor in the picture, a waveform monitor’s display can become distorted in some very strange ways. Mikko Wilson tested what happens with 2 cameras side-by-side.

Article.
When Barry Green tested Sony’s new EX1 camcorder, one of his tests was of course for resolution. The tired and true method of testing resolution involves using a waveform monitor to look at the actual signal in a very controlled way, allowing you to clearly and definitively see the changes from a test chart in the signal. In much the same way that a waveform monitor can be used to measure signal levels, it can also be used to measure timing - and therefore resolution - differences in a signal.

Barry’s initial test results showed a very high resolution from the EX1. However, as soon as he panned the camera, something caught him and others present off guard. Barry reported: “*Move [the camera], and you get a noticeable resolution drop, more than can be accounted for through simple motion blur. And it's not due to the codec either, you can see the effect on the live output and you can see it on a scope”* and later added “*everyone who watched the waveform monitor was like "whoa, what's happening here?"*”. Needless to say, this caused quite a heated reaction and much discussion as to whether the camera was actually loosing resolution as it was panned.

One theory, presented by member “TheMusician”, to explain what Barry had seen on the scope was the possibility that the skew caused by the CMOS sensor of the EX1 during panning would cause the details to skew across the one dimensional image of the waveform monitor which would resemble a notable resolution loss. Along with some theoretical test visuals, TheMusician’s theory sounded like a reasonable explanation as to what was happening, but as waveform monitors are something best analyzed in a real visual environment, I decided to put the theory to the test to see if I could replicate the phenomenon that Barry’s test had bought to light. TheMusician’s theoretical visuals formed a fantastic basis for a practical test.

If you are familiar with waveform monitors, skip this paragraph – If not, read on.
A waveform monitor is used to plot a video signal onto a graph with a scale so that it can be accurately measured. The x –axis is time, and the y –axis is the level of the signal (Hence “Y” being used to denote “Brightness” in video signals like Y/C [S-video] and YCbCr [Component]) with the width of the graph being equal to the width of one frame of video (the actual length of time varies depending on the video standard). Normally the signal information for the color information is filtered out, so a waveform monitor only displays the brightness of the black & white video signal. Unless a specific line is chosen for display, each and every line of the image is sequentially drawn on the scope, one on top of each other. Most scopes - especially digital “raster” scopes which also draw the whole frame at once - blank at the start of each frame. The result is that you see the brightness of every line in the image all at the same time, and this is where the problem presents itself. Every time the same level is present in multiple lines, the trace on the scope gets brighter as the “beam” re-draws over the same point. If every line of the image is identical, then the scope will have one single, bright, trace. This is the case for example with PAL color bars, which produce a

clean trace that looks like a flight of stairs, starting with the bright white bar on the left, and stepping all the way down to black at "0%" on the right.

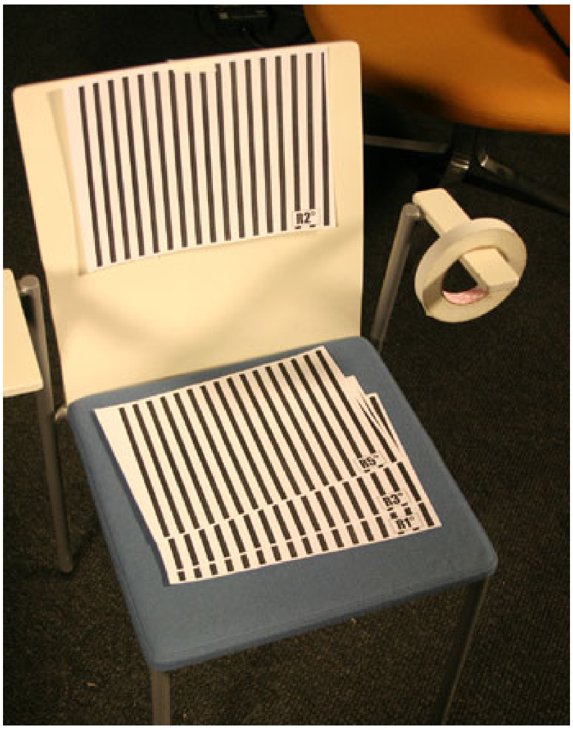

To test for the skew in the image from horizontal movement when captured with a CMOS sensor, the simplest test is one or more vertical bars. As color is irrelevant for this test, we'll use black & white bars.

The test chart setup can be seen on the right. The lines on the charts are skewed to the right by varying amounts to simulate/enhance the skew from the CMOS sensor. There is of course also a set of fully vertical bars as a "control" and the main subject of the test.

The test involved 2 cameras:
A Sony HVR-A1E, which has 1 1/3" CMOS sensor.
A Sony HVR-Z1E, which has 3 1/3" CCD sensors.

The imaging characteristics of both cameras are very similar, with the only major difference being the type of sensor used. The cameras where mounted on a tripod, stacked vertically above each other so that they would pan together at the same speed on the same axis. The only difference in the camera's vantage point is a sight difference in tilt-angle, but use these are vertical charts moving horizontally, tilt has a negligible effect on the test. The signal from each camera was then fed through a video mixer set to a vertical wipe, resulting in an image (& signal) that is half from a CMOS camera and half from a CCD camera. The cameras and full test setup, with accompanying video mixer, waveform/vector-scope rasterizser, and VCR can bee seen below.

The smaller HVR-A1 is on top, and the larger HVR-Z1 is below. The 'scope used was a Hamlet brand "MicroScope" which analyzes the incoming video signal and draws ("rasterizes") the scope graphs/plots onto the video signal itself. That combined signal can then be simply recorded onto a conventional VCR – in this case a standard MiniDV VCR. The MicroScope is the small blue unit under the VCR in the image below to the left. The video mixer is a Panasonic AW-SW350p with the frame-syncs turned on (as the cameras in the test can't be Genlocked). The combined output looks like the image on the lower right, which has some labels added in Photoshop – the small blue line denotes the location of the vertical wipe, which isn't apparent in the stationary control image.

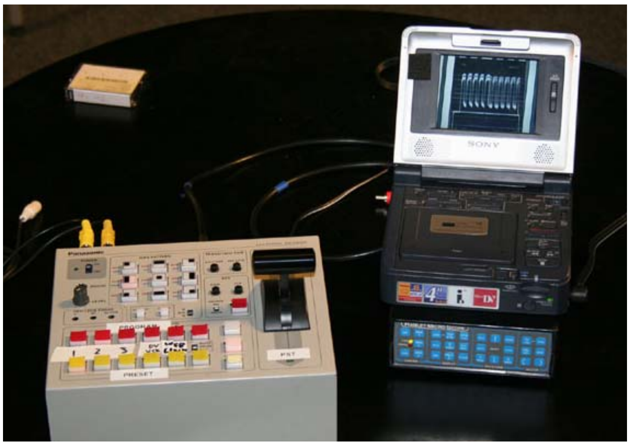

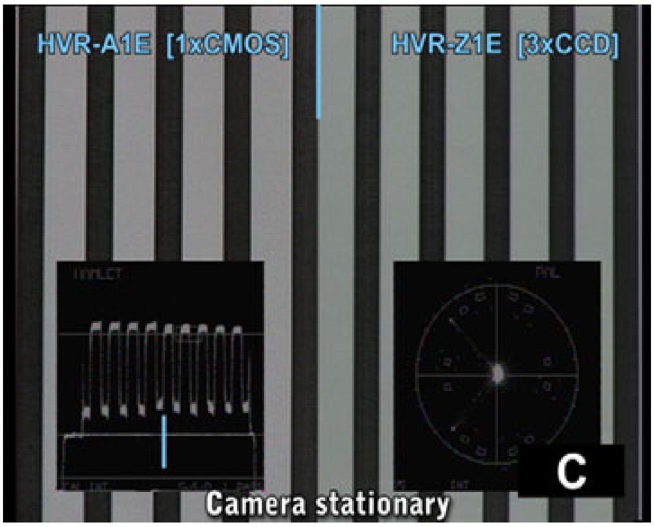

The image on the right is a static image (the cameras aren't panning) of the vertical "control" bars. Each horizontal line of video is identical to all the others, producing a nice clean single trace on the waveform monitor in the bottom left of the image. The lower points on the trace are the "Black" (actually dark gray) bars, and the higher points on the trace are the white spaces between the bars. Again, the light blue bar overlaid on the waveform display shows the separation between the signals from the two cameras.

So, we have our test setup: 2 similar cameras with differing sensor types carefully framed on the same test chart and able to pan together, images combined and analyzed together as one, and the critical split-screen waveform display. Now let's make it move…

The first test: panning the cameras over the control chart with vertical bars. Immediately the skew from the CMOS sensors is evident in the left side of the image and waveform plot, clear as day:

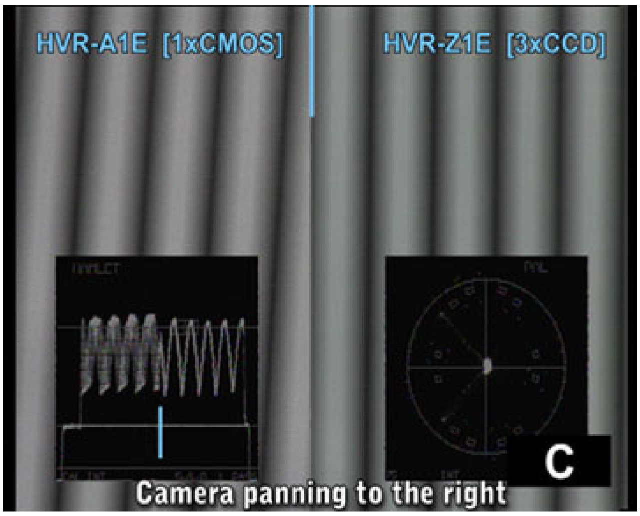

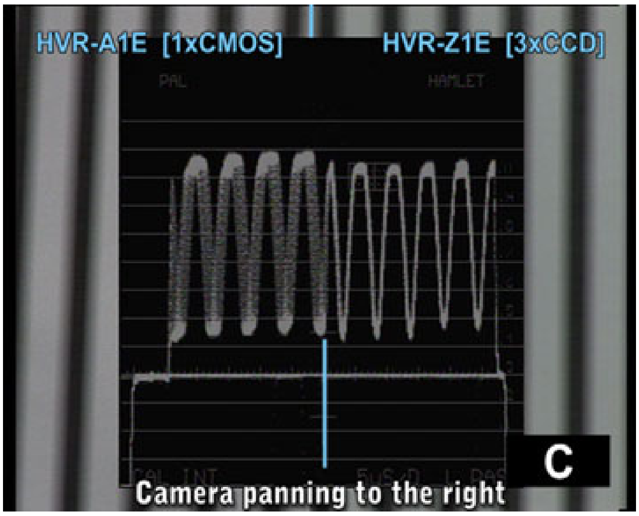

[Video clip “CWD_C.flv”]

As the camera pans to the right, the bars appear to move to the left. As the rolling CMOS shutter scans down the image, the bars move to the left as you look down the still image. The result: image skew. The global shutter of the CCD sensors exposes the entire image at the same time, and as a result, the bars appear vertical in the right side of the image. There is of course significant motion blur in the image, but as it’s only horizontal, it only affects the “sharpness” of the “peaks” and “valleys” of the plot on the waveform monitor. Both cameras were manually set to the same shutter speed, therefore the motion blur should be similar between them, and the visual image confirms this.

A quick technical note here about interlacing: both cameras are interlaced cameras. The still images were de-interlaced in Photoshop using the “Nearest neighbor” method, which simply removes one field (and half of the vertical resolution of the image) and has no effect on this test. Due to the interlaced output of the scope, this de-interlacing process also outputs the corresponding waveform display for the same lines as the visible image. In short: the cameras are interlaced, but they can be treated as half-resolution progressive for this test.

The result on the visible image is simple enough, but look what happens on the waveform, smaller in the left image, or larger on the right image (note that they are different pans, not necessarily at the same speed): The trace from the CCD camera is still nice and sharp as the bars are still vertical in the image and therefore each line in the image is the same – and on top of each other on the waveform monitor. But, and this is what Barry (and Co.) saw, the trace from the CMOS camera is heavily blurred, and initially indicative of a massive loss in resolution. What is actually happening is as the image skews slightly to the left with each successive line, each trace on the scope is shifted slightly to the left, resulting in a horizontally wider trace. The width of that trace depends not on the blur in the image – or its resolution – but rather simply on the amount of skew. The faster you pan, the ‘wider’ (“more blurred”) the trace on the scope is.

Even with a moderate pan, the result on simple single color test bars is considerable. It’s easy to imagine what this would do to a nice sharp trace of a more complex scene. For anyone looking purely at the scope image, or not noticing the skew in the source image (and it’s easy to miss if there aren’t any clear vertical bars in the image), would automatically be puzzled. And for anyone who’s use to looking at the waveform output of CCD cameras for years, the obvious initial assumed reason is exactly what Barry stated: “[This camera is loosing resolution when panning.]” But once we consider rolling shutter, the test simply becomes inconclusive as the results are distorted along with the image they are derived from.

As Adam and Jamie would say: Myth busted!

Then they would attempt to replicate the myth, or take things further. Although I didn't have anything to blow up as part of my tests, I also decided to explore further…

The theory that led to this test was taken another step in the online discussion. It was theorized that if rolling shutter caused vertical lines to skew when panned, would it not be possible to "straighten" lines that where already skewed on the test chart? And would this not then theoretically appear as an apparent sharpening of the image on the scope? As this sounded reasonable, I also tested for exactly that. Along with the "Control" chart of vertical black bars, I also printed a series of charts with the bars pre-skewed by a number of degrees. I repeated the above test on each chart, panning at various speeds manually trying to synchronize the pan speed with the rolling shutter so that the shutter would roll down the images the same rate as the skew of the bars.

And with the bars leaning just 1° to the right, the results were impressive. When panning to the right, the results where similar to the control but with the increase in lean apparent on both images. The CMOS camera's image was even more distorted and the distortion on the waveform more apparent to match. Also the CCD camera's output captured this lean of the bars, and began to distort their plot on the waveform too. This was pretty clear cut.

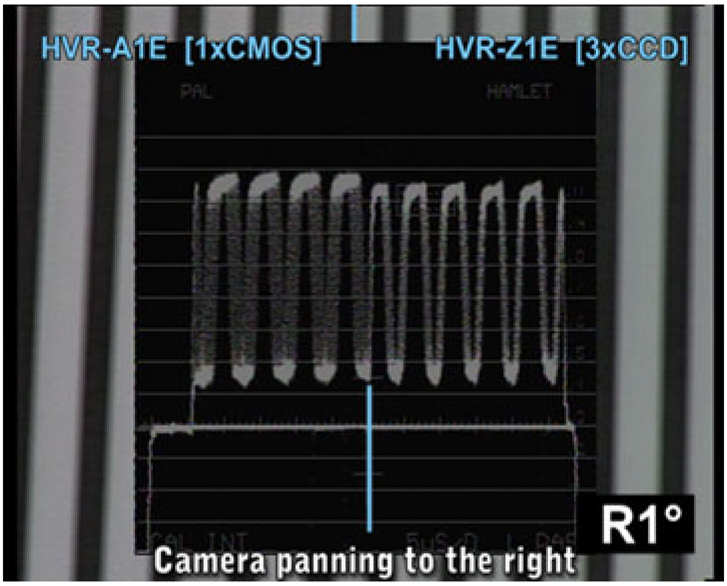

But watch what happened when the camera was panned to the left – opposite the direction of the bars.
There is no change to the output of the CCD camera: the bars still lean to the right, and the distortion on the waveform confirms it. But at the right panning speed, – and here's the cool part –the CMOS camera performed exactly as theorized: The rolling shutter skew "corrected" (counter distorted) the lean of the bars and made them appear vertical. And the result on the scope was most amazing: The image did in deed appear to sharpen during the pan!

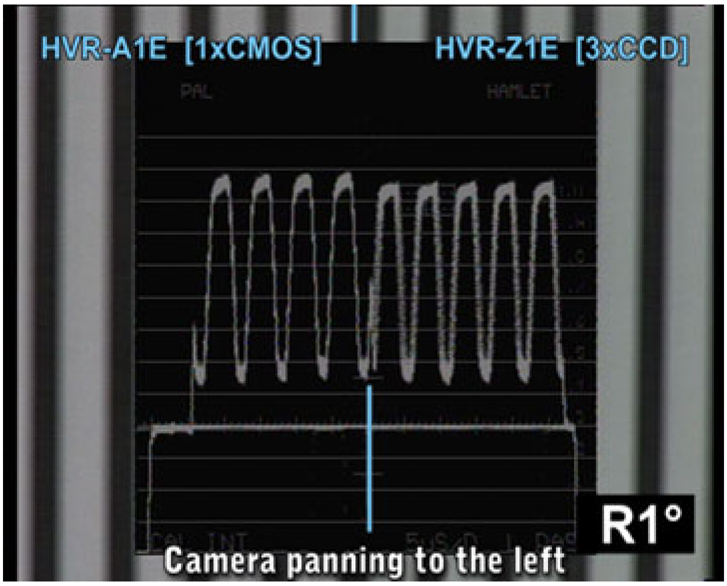

[Video clip “CWD_R1.flv”]

It was most amazing to watch the waveform plot stretching & compressing in real time in response to the incoming skewed image from the camera as the speed & direction of the pans varied.

Note that the CMOS image, though on first glace initially more “correct” with vertical bars, is still distorted. The bars in this image are *supposed* to be skewed. If this was a pan of a roofline of a building, this would look rather strange to have the roof “leaning” to one side.

As the test continued with increasing amounts of lean in the image, I had to pan the cameras faster and faster to keep the rolling shutter in sync with the bars. But the differences were more distinct on the CMOS image. (Remember that as global shutter CCD’s don’t skew, the panning speed had zero effect on that camera’s output.) As panning speed increased, so did motion blur; as a result, the test results become too blurred to interpret once the bars where leaning about 10°.

[Video clip “CWD_R10.flv”]

The most visibly extreme results I achieved where at about 3° of lean to the right. The images & video pretty much speak for themselves.

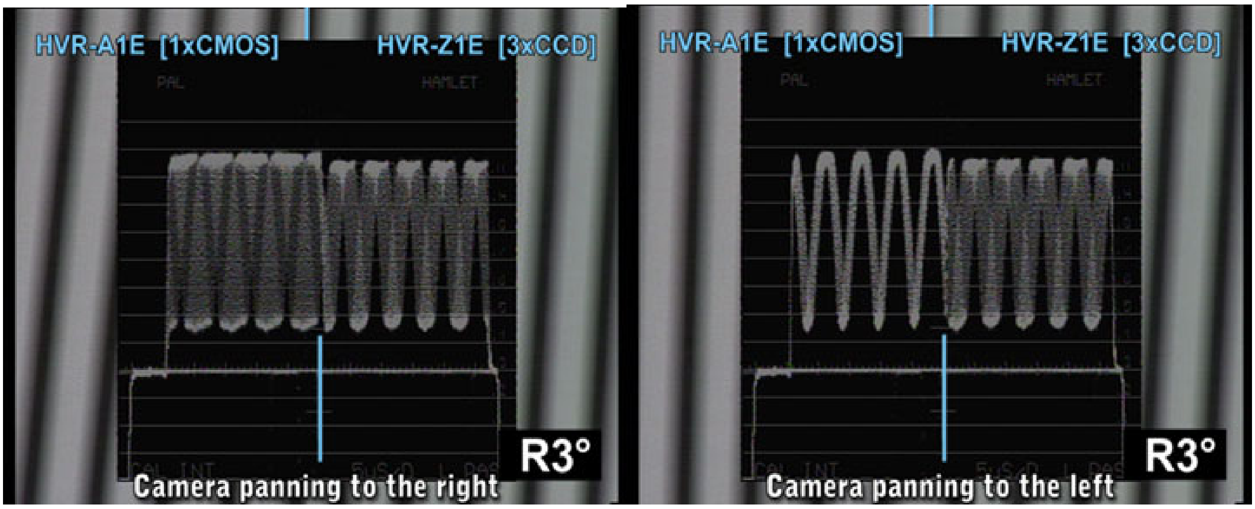

[Video clip “CWD_R3.flv”]

All of the test clips & images for a wider range of charts tested, along with larger setup photos, are available here: **http://dvxuser.mikkowilson.com/CWD/**

So in conclusion, visible skew from rolling shutter is directly proportional to the panning speed and the angle of any lines being panned across. Vertical lines and lines angled in the same direction of the pan will appear to lean more, and will appear more distorted on a waveform monitor. But also lines angle in the opposite direction to the pan will appear to “straighten” towards vertical, and their traces will become shaper on a waveform monitor.

With simple test patterns, these distortions are easy to notice and account for. But with real life images, the combination of distortion and counter-distortion can result in some very strange scope output. Add to that any vertical stretching or compression with vertical movement or tilting, and the waveform monitor’s image becomes just as subjective as the visual image. When measuring video, or anything for that matter, it’s important to keep in mind what you are measuring. As technology moves forwards, our methods to measure the results must also adapt, or we can find ourselves in trouble.

- Mikko Wilson

Mikko Wilson is a Media Engineer with 10 years of experience in both digital and analog video & TV production. He works as a freelance Steadicam Operator and engineering consultant. He has spent more time staring at scopes than he cares to think about, but enjoys exploring the advances and innovations in technology as part of an increasingly futile effort to thoroughly understand and keep up to date with the technology of our world today.