National Association of Broadcasters NAB TV TechCheck
THE WEEKLY NEWSLETTER FOR TV BROADCAST ENGINEERS

October 29, 2007

NEW DIGITAL VIDEO TEST PATTERN

At the SMPTE Technical Conference and Exhibition in New York last week, Sarnoff Corporation introduced a video test pattern for the visual evaluation of digital video systems. This new tool is a specialized video sequence that can reveal the quality and configuration of a video transmission, recording, and display chain, requiring only the use of a picture monitor and loudspeaker. The tests are easy to interpret and can be utilized for operational quality control in the broadcast environment as well as for system diagnostics and maintenance. System parameters covered include:

- Compression fidelity
- Lipsync
- Format conversion
- Colorspace mismatch
- Chroma subsamplng
- Frequency response
- Field reversal
- Highlight clipping
- Bit depth
- Skin tones
- Linearity
- Gamma

As shown in the figure, the test pattern comprises multiple sections combined into a single video image. Several of the video components are dynamic, where the changing display exhibits characteristics of the transmission path. The dynamic aspects of these tests are, of course, not apparent in the still illustration.

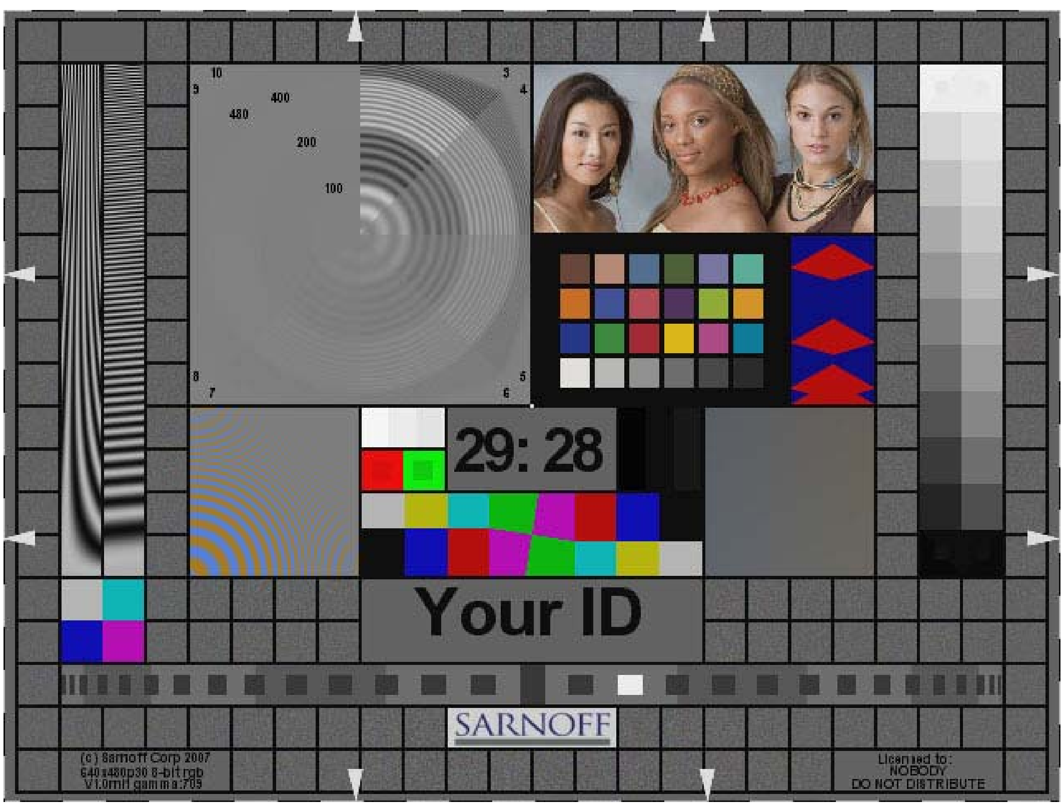

The test pattern is available in different versions for HD, SD and digital cinema resolutions and can be delivered as an A/V file on various digital media.

The following information on the various components of the test pattern relates to an interactive demonstration of the system provided to NAB by Norm Hurst of Sarnoff Corporation, who was responsible for the development of the test pattern.

Sweep

Purpose: To measure frequency response and aliasing horizontally and vertically; to measure relative performance of format conversion algorithms.

Description: The left-hand strip is a sweep of horizontal frequencies; the right-hand strip is a sweep of vertical frequencies. Each sweep starts at zero frequency at the bottom and increases to the Nyquist limit (alternating black-white-black-white-…) of the format at the top.

The sweeps are sinusoids with a minimum brightness value of black and a maximum of 75% of fullscale. They are angled slightly to accentuate aliasing that may occur, and they slowly change phase over time, which causes aliasing components to move in reverse.

Compression Fidelity

Purpose: To numerically assess compression fidelity.

Description: Radial sine wave whose frequency increases up to the Nyquist limit at the edges. The center represents 50 lines per picture height (LPH) and the frequency increases in evenly-spaced octaves (doublings) up to the Nyquist limit at the edges. A circular gray band indicates each octave and is labeled in LPH.

The pattern is divided into eight “pie segments”, each with different signal amplitude. The north-northeast (NNE) segment contains the largest-amplitude signal: the two most-significant bits (MSBs) are unchanging, so at least 3 bits of amplitude resolution are required to see detail. The next segment, east northeast, is one-half as large as the NNE segment, so the 3 MSBs are unchanging and it requires at least 4 bits to resolve detail. The “3 4” in the upper right corner indicates that these segments correspond to bit depths of 3 and 4 bits.

Each segment is one-half the amplitude of the previous segment and requires one more bit (half the quantization step size) than the previous segment in order to resolve detail. The pairs of numbers in the corners indicate the bit depth required to resolve detail for the two segments in that quadrant. The bit-9 and bit-10 segments have no detail in this example because an 8-bit image file was used to generate this example.

Usage: Shows the effect of compression at various quality settings. Look for regions where the detail vanishes, indicating that the quantization step size is larger than the amplitude of the detail in that region. For each bit-slice, note which octaves retain detail and which do not.

Chroma Zone Plate

Purpose: To indicate chroma resolution, e.g. 4:2:2, 4:2:0.

Description: Radial sine wave whose frequency increases up to the Nyquist limit at the edges. The lower left corner represents 50 lines per picture height and the frequency increases in evenly-spaced octaves (doublings) up to the Nyquist limit at the edges. The amplitude is constant. The figure at the right shows full-resolution, 4:4:4 chroma. The luminance is constant, and remains constant when transcoded to either ITU-601 or ITU-709 colorspace.

Usage: Select from the pulldown menu to see subsampled versions with and without prefiltering.

Skin Tone Reference

Purpose: To provide a quick, subjective check of image fidelity.

Description: Professionally lit and photographed models representing a range of skin tones.

Usage: When it comes to faces, the human visual system is very sensitive—small errors register with our brains very quickly. Check for accurate skin tone color and gradations. Check for detail such as hair in dark and light areas. Check that highlights are not blown out, and check that dark areas are not blocked up.

SMPTE 303M

Purpose: To check proper rendering of colors.

Description: A 6 x 4 array of color patches, whose color values are as set forth in SMPTE 303M "Color Reference Pattern" and gamma-encoded as noted in the Format feature.

The first row contains colors found in nature: Dark skin, Light skin, Blue sky, Foliage, Blue flower, Bluish green. The second row also contains natural colors: Orange, Purplish blue, Moderate red, Purple, Yellow, green, Orange yellow. The third row contains additive and subtractive primaries: Blue, Green, Red, Yellow, Magenta, Cyan. The last row consists of a six-step neutral gray scale; the neutral values are Munsell values (the R, G, and B values are not identical).

Usage: Experts familiar with this chart may evaluate it against what they expect to see. Otherwise, computer software may be used to objectively measure the color values of each square and compare those against the values specified in SMPTE-303M.

Lava Lamp

Purpose: To check progressive scan conversion and 4:2:0 interlaced subsampling, and to check for smooth motion rendition.

Description: Accelerating red diamonds over blue background. The red is 75% red, and the blue is 50%. The speed at the ¼ way point is one line per frame, or one line per field for interlaced formats. The maximum speed is 2 lines per frame, or 2 lines per field for interlaced formats. The diamonds are spatially anti-aliased to mimic optically smooth motion.

The slopes on the upper half of each diamond are 16 lines from side to center. In the lower half, the slopes are 8 lines from side to center. Note that the absolute slope, or height-to-width ratio of the diamonds, depends on the vertical resolution of the particular format. The number and spacing of the diamonds depends on the frame rate or field rate of the particular format.

Usage: Progressive scan conversion: Note the "jaggies" along the edges of the diamonds. Note how the jaggy appearance changes as the speed increases. Better progressive scan converters will maintain more resolution at higher speeds.

4:2:0 interlaced chroma: Systems like MPEG that vertically subsample the chroma of interlaced video (i.e. 4:2:0 chroma) specify that the chroma from separate fields be kept separate, so that the motion of the luma and the chroma match. An error will be seen as the red portion juddering while the gray portion moves smoothly. Also, errors in encoding or reconstruction of interlaced 4:2:0 chroma may result in the diamonds having very rough edges.

Smooth motion: disturbances in the smoothness of the motion are easily seen.

White Pluge

Purpose: To set gain (contrast) of a display, especially for LCD displays, to verify that highlights are not blown out and to check that encoders and decoders do not clip off the signal's headroom. It is also used in conjunction with the Colorspace Mismatch Indicator.

Description: Three white strips: 104%, 100%, and 96% of full scale.

The sweeps are sinusoids with a minimum brightness value of black and a maximum of 75% of fullscale. They are angled slightly to accentuate aliasing that may occur, and they slowly change phase over time, which causes aliasing components to move in reverse.

Usage: Before using, adjust display offset (brightness) using the Pluge feature. Reduce display gain (contrast) until all three strips are visible. (If the boundary between the left and middle strips cannot be discerned even after reducing the contrast, check upstream processes to see if one of them clipped off the headroom.)

For LCD displays, increase the gain until the boundary between the left and middle strips cannot be seen. (For CRT displays, increase the gain as desired, but not beyond the point where the first and second strips merge.)

Colorspace Mismatch Indicator

Purpose: To indicate when clipping occurs as a result of decoding YCbCr with the incorrect color matrix.

Description: Red and green squares with 100% border and a reduced-level inner patch.

Usage: First use the White Pluge feature to force the display to clip at 100%: increase the contrast until the boundary between the two brightest strips just disappears (LCDs work well for this). Then verify that the inner patches are visible in the Color Mismatch Indicator. If the red patch blends in with the red square, this may indicate that YPbPr coded as ITU-601 was decoded using an ITU-709 matrix. However, if the green patch blends in with the green square, this may indicate that YPbPr coded as ITU-709 was decoded using an ITU-601 matrix.

Pluge

Purpose: To set black level (brightness) of a display; to check that encoders and decoders do not clip off the signal's footroom.

Description: Three dark strips: -4%, 0%, +4%, as found in standard color bars conforming to SMPTE Engineering Guideline EG 1-1990.

Usage: Adjust display offset (brightness) until all three strips are visible, then reduce the brightness until the boundary between the left and middle strips cannot be seen. If the display cannot be adjusted to show three separate strips, check upstream equipment for something that may have clipped off the footroom.

Countdown

Purpose: To provide a countdown for tape leaders; to label each frame and field of the test sequence with a unique number; to detect frame drops and freeze-frames; to measure large lipsync errors.

Description: Black on gray numbers indicating seconds and frames remaining to the end of the sequence. The final frame says “00:01”. For interlaced formats, the colon is replaced with a semi-colon for the first field of each frame. The largest value for the frame count depends on the frame rate (e.g. 29 for 30 fps, 23 for 24 fps).

Usage: Use as a leader countdown. Use in frame-by-frame analysis to keep track of which frames have which artifacts. Measure large lipsync errors (greater than one second) using wide-spaced events in the audio track.

Color Bars

Purpose: To check for correct levels and offset, and to check chroma subsampling phase.

Description: Standard 75% split-field bars with a tilted transition between green and magenta. The tilt is off-vertical by a slope of 8 pixels over the height of the bars. It is off-horizontal by 8 lines over the width of the green and magenta bars.

To highlight edge behavior, the chroma in the 4:2:2 rendering of the color bars is subsampled without prefiltering.

Usage: The tilted bars may be used to examine the performance of chroma upsampling, and to verify correct phasing of the interpolation with respect to odd and even pixels and lines.

Shallow Ramps

Purpose: To measure visibility of contouring due to quantization of luma or chroma.

Description: In the first frame there is a shallow, two-dimensional luma ramp rising 12/255 of fullscale from the upper left corner to the lower right corner, and a shallow blue ramp, perpendicular to the luma ramp, rising 24/255 of fullscale from the lower left corner to the upper right corner. The two ramps rotate clockwise together, completing one revolution every two seconds. The figure shows an 8-bit rendering; the diagonal contour lines are clearly visible.

Usage: Look for rotating contour lines, indicating 8-bit quantization. Conversion from 10 bit to 8 bit video with careful use of dithering may render the contour lines invisible. This feature is particularly difficult to encode well.

Stairstep

Purpose: To measure gamma, linearity of luma, and to set display range.

Description: Eleven luma values evenly distributed over the 0 to 100% range. The left side is linear “voltage”; the right side is linear light for a particular gamma. The gamma function used to render the pattern is indicated in the “Format” box. For most formats the gamma function is the one specified by ITU-709.

The difference between adjacent steps is 10% of full range.

The white step at the top and the black step at the bottom have the same value for both linear voltage and linear light. These steps have additional features, shown in detail below.

The black and white steps each have diagonal “dog ears” that exceed the maximum limits by 2% on the left and 1% on the right. (All percentages are relative to linear voltage.) They also contain small squares that are within the limits by 2% and 1%. These features blink on and off. The 2% features blink once per second; the 1% features blink at half that rate.

Usage: The four, small blinking squares should be visible, and the dog ears in black should not be visible. LCD monitors can be adjusted to make the white dog ears vanish as well.

Adjust monitor brightness until the 1% black dog ears are not visible but the 1% square remains visible. For LCD-type displays, adjust monitor contrast until the 1% white dog ears vanish but the 1% square remains visible. Check the brightness adjustment and repeat as necessary.

Lipsync

Purpose: To measure audio/video delay mismatch; to measure chroma motion judder; to detect interlaced field reversal.

Description: Moving mark over a scale of tic marks, and an associated audio track. The mark is 100% white when moving left-to-right and 100% red when moving right-to-left. When properly synchronized, the audio beep will be heard when the white moving mark appears at the center. Alternating light gray and dark gray bars indicate 100 ms each. There is one scale tic mark per frame or field; the number of tic marks in each 100 ms bar depends on the frame rate at which the pattern was created. Interlaced formats have a tic mark for each field.

Usage: Rather than follow the moving mark with your eye, choose one stationary mark and watch for it to blink while listening for the beep in the audio. If the sound is heard before the flash, choose a different tic mark to the left and watch again. If the sound is heard after the flash, choose a different tic mark to the right and watch again. Repeat until you find the tic mark that seems best synchronized with the sound. Count the number of tic marks from that tic to the center to determine the number of frames (or fields for interlace) of delay offset. With practice, lipsync error can be measured to 1 frame accuracy.

Interlaced video with chroma processed as non-interlaced will seem to show not one but three red squares, with the middle one brighter than the other two (see below). This may be more apparent on a CRT monitor compared to an LCD monitor.

When shown using a display that incorporates a color wheel to achieve RGB, will often exhibit color-fringing on the left and right sides of the moving white mark.

Audio

Purpose: To measure audio/video delay mismatch; to determine left/right channel swap; to measure audio levels.

Description: Two-channel audio track with ticks and continuous tone. One tick appears in both channels every 2.000 seconds (or 2.002 seconds for formats running 1000/1001 slow). This tick occurs each time the Lipsync indicator flashes at the middle point. Another tick is heard in the right channel when the Lipsync indicator reaches the right side, and a similar tick is heard when the indicator reaches the left side.

A continuous tone is heard in both channels. This tone has a level of –20 dB FS. The tone mutes for 2 seconds every 10 seconds. Just after the tone mutes, the blinker feature blinks white instead of black.

Usage: See the description of the Lipsync and Blinker features to use for lipsync measurements. Watch and listen to the left and right ticks to determine if the left and right channels are reversed: they should correspond spatially with the movement of the Lipsync indicator. Note the level on a VU meter and adjust gain until it reads –20 dB FS.

Lipsync Blinker

Purpose: To measure audio/video delay mismatch using a photodetector and an oscilloscope.

Description: Black or white rectangle that turns to gray when the lip sync audio track beeps. The duty cycle ratio is 1.5 seconds of gray, 0.5 seconds of black or white. Black appears except when the seconds digit of the countdown is “2” white appears.

Usage: To measure gross a/v sync, listen for the 1000 Hz tone to stop, then observe the blinker. It should blink white just after the tone stops. If not, a/v sync may be very far off.

To make accurate measurements of small A/V errors: place a photodetector over the blinker and use an oscilloscope to look at its output. Also look at the audio output with the scope simultaneously. A storage scope might make this easier. Compare the timing of the black-to-gray and white-to-gray transitions of the blinker to the time of the beep in the audio track.

This method can be used to measure audio/video timing errors to better than 5 ms accuracy.

Blue Only

Purpose: To set chroma gain (“color”) and phase (“hue” or “tint) of NTSC displays.

Description: White, cyan, blue and magenta chips, 75% of fullscale. The white is over the blue and the cyan is over the magenta. This mimics the arrangement of the “reversed bars” of SMPTE Engineering Guideline EG 1-1990.

Usage: If the monitor has a “blue only” button, press it. Alternatively you can view the image through a blue filter, such as the Kodak Wratten Blue 47B Gelatin filter or the Lee Colortran Congo Blue #181.

Adjust the chroma gain and phase until all four chips have the same intensity and appear to blend together.

For information on Sarnoff compliance bitstreams, see: http://www.sarnoff.com/bitstreams. Norm Hurst may be contacted at: nhurst@sarnoff.com.

The Shift from Daylight Saving Time Nears

Reminder: All DTV broadcasts should now be announcing in their PSIP that daylight saving time will end on **November 4, 2007**.

The ATSC Standard A/65 requires that one month before the shift from daylight saving time that data announcing the shift be sent. As A/65 is part of the FCC regulations, this is required for all DTV transmissions. Specifically; the following are the proper settings for the following STT fields (as of 2 AM, October 4, 2007):

DS_Status = 1

DS_Day_of_month=4

DS_Hour=2