XAPP682 (v1.0) March 6, 2004

HDTV Video Pattern Generator

Author: John F. Snow

Summary

This application note describes a technique for generating high-definition (HD) video test patterns with Xilinx FPGAs. Video test patterns are used to verify the proper operation of video equipment. Most video equipment capable of generating a video signal can also produce one or more video test patterns to verify proper operation of the video equipment.

The video test pattern generator presented here uses the 18 kb block SelectRAM memory present in the Virtex-II™, Virtex-II Pro™, and Spartan™-3 FPGA families from Xilinx. The block memories are used to hold the video patterns. Several different video patterns can be stored in the block memories. The video pattern generator described in this application note produces three commonly used HD video test patterns. The video test pattern generator supports all 18 of the HD digital component video formats supported by the HD-SDI (SMPTE 292M) standard. The video test pattern generator uses very few FPGA resources, and it can easily be placed in the same device with other video processing functions.

HD Digital Component Video

Professional broadcast studios and video production centers typically use digital component video as the preferred video format for content creation, storage, and editing. Xilinx application note XAPP248, "Digital Video Test Pattern Generators", contains a description of digital component video and describes a video pattern generator for standard-definition video. [Ref 1]

Video Format Naming

There are two basic scanning methods used for video: interlaced and progressive. The industry standard naming scheme for video formats uses the number of active lines followed by either "i" for interlaced or "p" for progressive. An interlaced video format with 1080 active lines would be called 1080i. Often, the frame rate is also added to the end of the video format name. For example, a 1080i format with a 30 Hz frame rate would be called 1080i30.

In interlaced scanning, the frame is split into two fields with one field containing the odd lines and the other frame containing the even lines. When displayed on a monitor, all the lines from one field are drawn, followed by all the lines from the next field. The two fields are actually from different images taken at different times. There are two interlaced video formats currently defined for HDTV: 1035i with 1035 active lines, and 1080i with 1080 active lines. The 1035i format was an early HDTV format with non-square samples and has been almost entirely replaced by the 1080i formats.

In the progressive formats, the frame is not split into fields. All of the lines of the frame are from the same image. When displayed on a monitor, all of the lines of the frame are displayed sequentially from top to bottom. There are two progressive video formats currently defined for HDTV: 720p with 720 active lines and 1080p with 1080 active lines.

SMPTE recommended practice RP 211 adds a third scanning method called segmented frame. [Ref 2] Segmented frame scanning is really progressive video that has been reformatted to make it compatible with interlaced equipment. In the early days of HDTV development, equipment for progressive scan video, particularly video recorders, was not as readily available

as equipment for interlaced video. The segmented frame technique provided a standard method for taking progressive scan video and reformatting it to be compatible with video recorders and other equipment designed for interlaced video. In segmented frame video, each progressive scan frame is separated into two fields, one containing the even lines and the other containing the odd lines. Thus, the video appears to be interlaced. It is, in fact, indistinguishable from interlaced video. However, the two fields in the segmented frame represent one progressive frame and are taken from one image whereas, the two fields of an interlaced frame represent separate images taken at different times. Before being displayed, the two fields of a segmented frame are recombined into a progressive frame for display on a progressive scan monitor. The segmented frame formats use the letters "sF" in the format name like this: 1080sF30.

Digital Component Video Standards

The most commonly used digital component video formats used in the broadcast industry today are based on the 4:2:2 sampling scheme using the Y'Cb'Cr' color space. Table 1 shows a list of commonly used HD 4:2:2 component digital video standards. All of these formats are compatible with the SMPTE 292M HD-SDI standard for transporting HD digital video over coax cable or optical fiber. The last column shows the SMPTE 292M format designations for these video standards. The format designation for the segmented frame formats show the equivalent 1080p standard.

Table 1: **Common HD 4:2:2 Component Digital Video Standards**

SMPTE Standard	Format	Frame Rate (Hz)	Sample Rate (MHz)	Active Samples/Line & Active Lines/Frame (words x lines)	Total Samples/Line & Total Lines/Frame (words x lines)	Format Designation
SMPTE 260M	1035i	30	74.25	1920 x 1035	2200 x 1125	A
SMPTE 260M	1035i	30/M	74.25/M	1920 x 1035	2200 x 1125	B
SMPTE 295M	1080i	25	74.25	1920 x 1080	2376 x 1250	C
SMPTE 274M	1080i	30	74.25	1920 x 1080	2200 x 1125	D
SMPTE 274M	1080i	30/M	74.25/M	1920 x 1080	2200 x 1125	E
SMPTE 274M	1080i	25	74.25	1920 x 1080	2640 x 1125	F
SMPTE 274M	1080p	30	74.25	1920 x 1080	2200 x 1125	G
SMPTE 274M	1080p	30/M	74.25/M	1920 x 1080	2200 x 1125	H
SMPTE 274M	1080p	25	74.25	1920 x 1080	2640 x 1125	I
SMPTE 274M	1080p	24	74.25	1920 x 1080	2750 x 1125	J
SMPTE 274M	1080p	24/M	74.25/M	1920 x 1080	2750 x 1125	K
SMPTE 296M	720p	60	74.25	1280 x 720	1650 x 750	L
SMPTE 296M	720p	60/M	74.25/M	1280 x 720	1650 x 750	M
SMPTE RP 211	1080sF	30	74.25	1920 x 1080	2200 x 1125	(G)
SMPTE RP 211	1080sF	30/M	74.25/M	1920 x 1080	2200 x 1125	(H)
SMPTE RP 211	1080sF	25	74.25	1920 x 1080	2640 x 1125	(I)
SMPTE RP 211	1080sF	24	74.25	1920 x 1080	2750 x 1125	(J)
SMPTE RP 211	1080sF	24/M	74.25/M	1920 x 1080	2750 x 1125	(K)

Color Space and 4:2:2 Sampling

Monochrome TV signals contain only intensity information, called luminance, or luma, and designated with the letter Y' (the apostrophe indicates that the component is gamma corrected). When color information was added to the TV signal, the luma signal was left intact for compatibility with existing equipment, and two components of color information, called U' and V', were added. The two color components are often called color difference signals. The U' component is the difference between blue and luma. The V' component is the difference between red and luma.

Digital component video uses a modified form of the Y' U' V' color space. This color space, called Y' Cb' Cr', is a scaled and offset version of Y' U' V'. The Cb' component is the blue color difference component, similar to the U' component in Y' U' V'. Likewise, Cr' is the red color difference, similar to V'. In 10-bit digital video, the Y' component has a range of 64 to 940. The Cr' and Cb' components have ranges of 64 to 960. Values above and below the specified component ranges are reserved.

In 4:2:2 digital component video, the Y' component is sampled at twice the rate as each of the chroma components Cb' and Cr'. For example, if the basic sample rate is 74.25 MHz, then the Y' component is sampled at a rate of 74.25 MHz. The Cb' and Cr' components are sampled at half this rate, or 37.125 MHz. Because there are two chroma components that are each sampled at half the sample rate, there is an average of one chroma word per sample. So, each sample of video consists of one word of Y' and one chroma word, either Cb' or Cr'. Consecutive samples alternate between having a Cb' word and a Cr' word.

In the HDTV digital component video standards, the luma channel and the chroma channel are treated separately. For 10-bit digital video, there is a 10-bit luma channel and a 10-bit chroma channel. In Figure 1, notice how consecutive samples alternate between having a Cb' word and a Cr' word on the chroma channel.

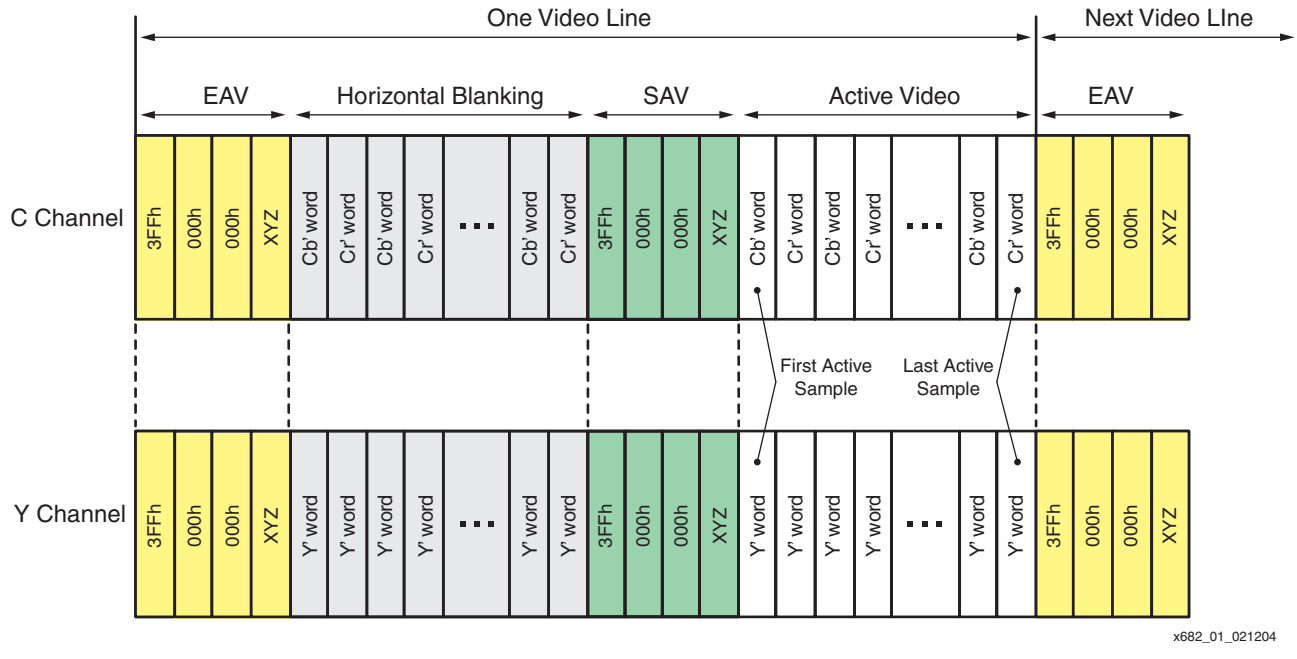

Figure 1: **HDTV Video Line Format**

Timing Reference Signals

Each line of HD digital component video is formatted as shown in Figure 1. A line is divided into a horizontal blanking interval and the active video portion. The end-of-active video (EAV) and start-of-active video (SAV) timing reference signals serve to mark the transitions between the active and blanking portions of the line. Each EAV and SAV is four words long. The first three

words are always 3FFh, 000h, and 000h in that order. The fourth word is called the XYZ word and contains three timing flags used to indicate which field the line is in, whether the line is an active video line or is in the vertical blanking interval, and whether this is an EAV or an SAV. Figure 2 shows the format of the XYZ word.

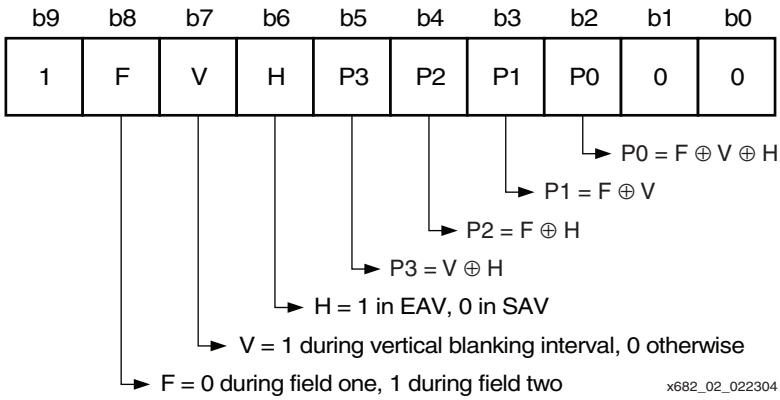

Figure 2: **XYZ Word Format**

The timing reference signals play a key role in the HD-SDI protocol. The leading sequence of 3FFh, 000h, and 000h is unique in the video stream. This sequence can only occur legally at the beginning of a timing reference. HD-SDI receivers use this sequence to synchronize to the HD-SDI bit stream.

Video Test Patterns

Many different video test patterns have been developed to test different aspects of video transmission, reception, and display. This application note focuses on three video test patterns commonly used in the broadcast studio, often in conjunction with the HD-SDI standard. The three test patterns are: SMPTE RP 219-2002 color bars, 75% color bars, and the SMPTE RP 198-1998 HD-SDI checkfield.

SMPTE RP 219-2002 Color Bar Pattern

SMPTE recommended practice RP 219-2002 specifies a color bar pattern that is compatible with both standard- and high-definition equipment. It is similar to the older standard-definition SMPTE EG-1 color bar pattern, but has been updated to include a 16:9 aspect ratio pattern and some new features. Figure 3 shows this test pattern.

The SMPTE RP 219-2002 document defines the sizes of the various areas in the color bar pattern and also explicitly specifies the digital luma and chroma values for each of the colors. Consult this document for more information about the pattern.

One of the features of this color bar pattern is the Y-Ramp pattern. This pattern ramps linearly from 0% luma (black) to 100% luma (white) from left to right across the pattern.

Figure 3 shows the 16:9 aspect ratio version of this pattern. The 4:3 aspect ratio version leaves off the two columns (40% gray in the top portion) on either side of the pattern.

75% Color Bars Pattern

The top row of SMPTE RP 219-2002 is a 75% color bar pattern. Sometimes it is useful to remove the lower portions of the SMPTE RP 219-2002 pattern and have the 75% color bars pattern occupy the full height of the frame. In particular, the SMPTE 292M HD-SDI spec requires that the output jitter of an HD-SDI transmitter should only be measured with 75% color bars. Figure 4 shows this test pattern.

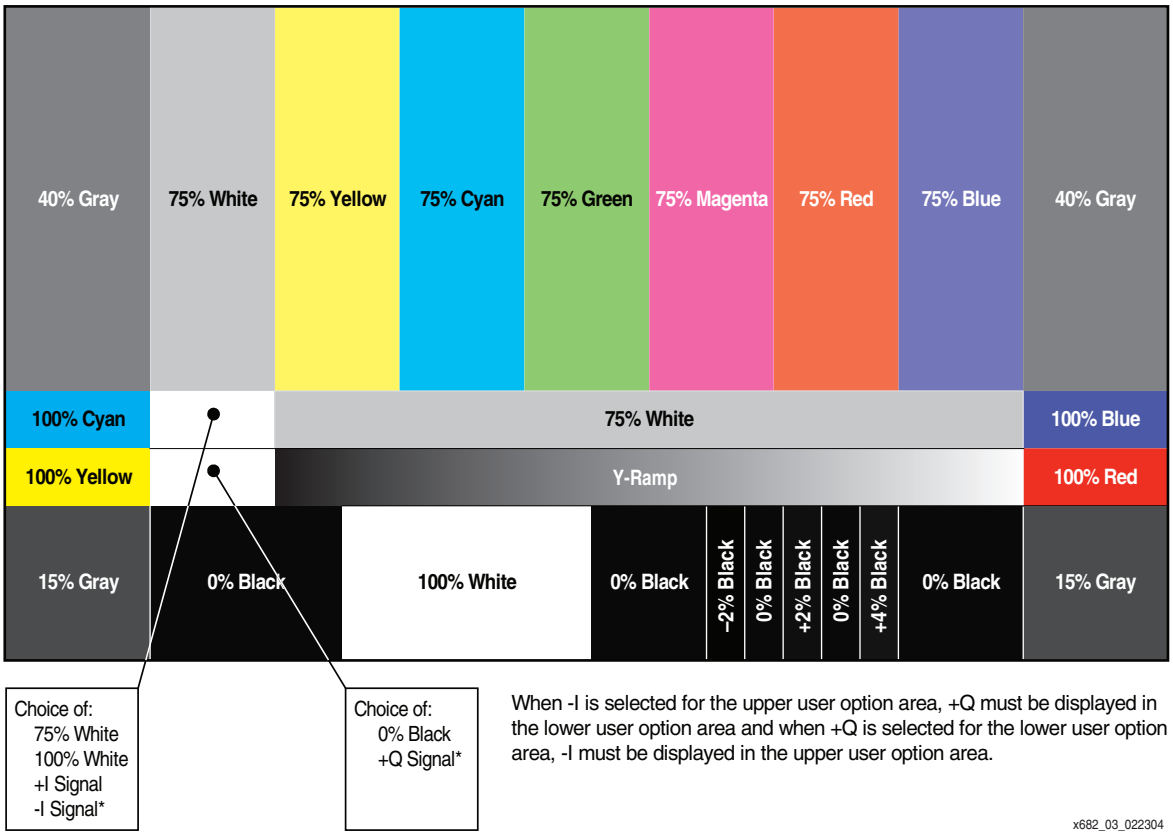

Figure 3: **SMPTE RP 219-2002 Color Bar Pattern (16:9 Aspect Ratio)**

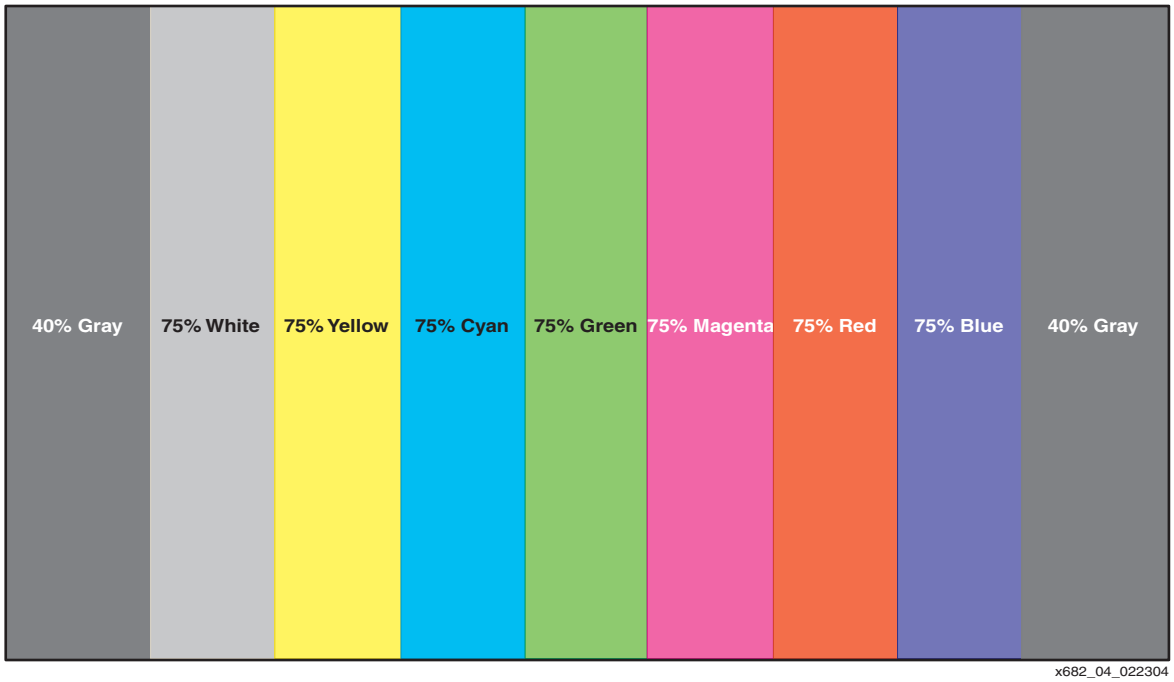

Figure 4: **75% Color Bars Pattern (16:9 Aspect Ratio)**

HD-SDI Digital Checkfield Pattern

SMPTE recommended practice RP 198-1998 defines a special pattern designed to generate the HD-SDI pathological waveforms. There are two different pathological waveforms. One pattern stresses the HD-SDI receiver's cable equalizer and the other stresses the receiver's PLL. The two pathological waveforms are shown in Figure 5.

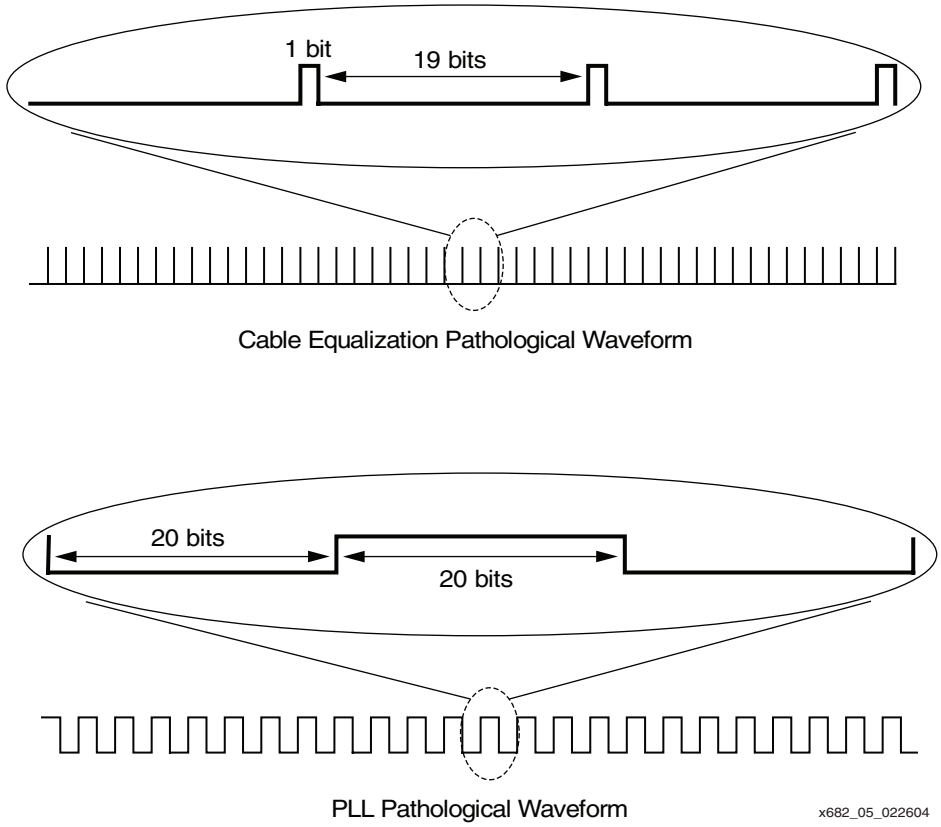

Figure 5: **HD-SDI Pathological Waveforms**

The equalizer pathological waveform is poorly DC balanced. It consists of a single High bit followed by 19 Low bits. This basic pattern is repeated continuously across a video line. Note that the opposite polarity of this waveform, a single Low bit followed by 19 High bits, is also possible and both polarities are generated by the RP 198-1998 checkfield.

The PLL pathological waveform is a square wave with 20 Low bits followed by 20 High bits. This basic pattern is repeated continuously across a video line. This pattern is much lower in frequency than the typical bit patterns seen by the receiver's PLL. Poorly designed PLLs may not stay locked during this low frequency waveform.

It is important to distinguish between the RP 198-1998 checkfield and the pathological waveforms. The pathological waveforms are serial bit patterns that are output by an HD-SDI transmitter after encoding and serialization. The checkfield is a video pattern that will, when encoded by an HD-SDI transmitter, produce the pathological waveforms.

The HD-SDI encoding process is based on a pseudo-random scrambling algorithm. The encoding of each sample varies depending on the current state of the scrambler. The state of the scrambler is dependent upon the previous history of the data that has been encoded by the scrambler. Therefore, it is not possible to force the HD-SDI encoder to generate either pathological waveform "on demand". The RP 198-1998 checkfield causes an HD-SDI encoder to produce the pathological waveforms randomly with a certain statistical rate of occurrence.

The RP 198-1998 checkfield is shown in Figure 6. In the top half of each field or frame, the checkfield consists of repeatedly sending a luma value of 300h and a chroma value of 198h into the encoder to cause the it to occasionally generate the equalizer pathological waveform. These values can be swapped with the chroma channel set to 300h and the luma channel set to 198h; the order doesn't matter. While encoding half a field or frame of this pattern, the equalizer pathological pattern will usually be produced on several video lines in the frame.

In order to insure that both polarities of the equalizer pathological waveform are produced by the encoder, the 198h value in the first sample of the first active picture line in the first field of every other frame is changed to a value of 190h. For progressive scan formats, the sample that is changed is the first sample of the first active line of every other frame.

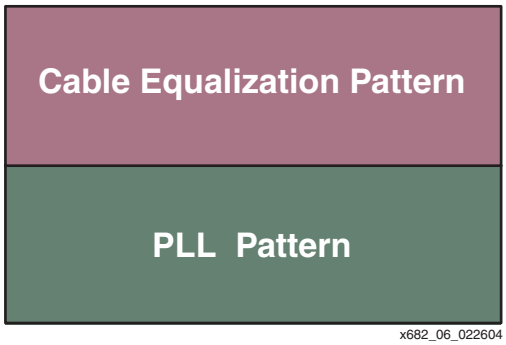

Figure 6: **RP 198-1998 Checkfield**

The bottom half of each field or frame in the checkfield consists of the pattern 200h, 110h. As with the equalizer pattern, it doesn't matter which value is assigned to which channel. This sequence causes the HD-SDI encoder to produce the PLL pathological waveform during several lines of the frame (on average).

Reference Design

Introduction

The HD video pattern generator reference design can produce three different video patterns: SMPTE RP 219-2002 color bars, 75% color bars, and the SMPTE RP 198-1998 checkfield. In all cases, the pattern generator produces patterns with 16:9 aspect ratios.

The reference design can produce the video patterns for all of the video formats supported by HD-SDI (all of the formats listed in Table 1). The video pattern generator actually only has to produce 9 different video formats as listed in Table 2 because some of these video formats can be combined. For example, the only difference between format D (1080i 30 Hz) and E (1080i 30/M Hz) is the clock rate. The pattern generator produces exactly the same data for both of these formats and the frequency of the video clock must be switched externally to the video pattern generator. Similarly, with a static image, such as these video test patterns, the data that the video pattern generator produces for the "look-alike" segmented frame and interlaced formats is identical. So, for example, the video pattern generator produces the same data for both the 1080i 30 Hz format and the 1080sF 30 Hz format.

The video pattern generator reference design can produce all nine of the video formats listed in Table 2. However, it can only be loaded with eight of these formats at any one time. By default, the reference design does not include the SMPTE 260M 1035i 30 Hz video format (group 0 from Table 2), since it is rarely used. However, it is a simple matter to change the reference design to include any set of eight video formats shown in Table 2.

Table 2: **Video Format Groups Supported by XAPP682 Video Pattern Generator**

Group Number	Supported Video Formats (SMPTE 292M Format Designator Shown in Parenthesis)			
	Progressive or Interlaced Formats		Segmented Frame Formats	
	74.25 MHz	74.25/M MHz	74.25 MHz	74.25/M MHz
0	1035i 30 Hz (A)	1035i 30/M Hz (B)		
1	1080i 25 Hz (C)			
2	1080i 30 Hz (D)	1080i 30/M Hz (E)	1080sF 30 Hz	1080sF 30/M Hz
3	1080i 25 Hz (F)		1080sF 25 Hz	
4	1080p 30 Hz (G)	1080p 30/M Hz (H)		
5	1080p 25 Hz (I)			
6	1080p 24 Hz (J)	1080p 24/M Hz (K)		
7	720p 60 Hz (L)	720p 60/M Hz (M)		
8			1080sF 24 Hz	1080sF 24/M Hz

Figure 7 is a block diagram of the reference design video pattern generator. The top level design files for this reference design are named `multigenHD.v` and `multigenHD.vhd`. This design is built around three 18 kb block RAMs. The block RAMs are used as ROMs and are initialized during FPGA configuration.

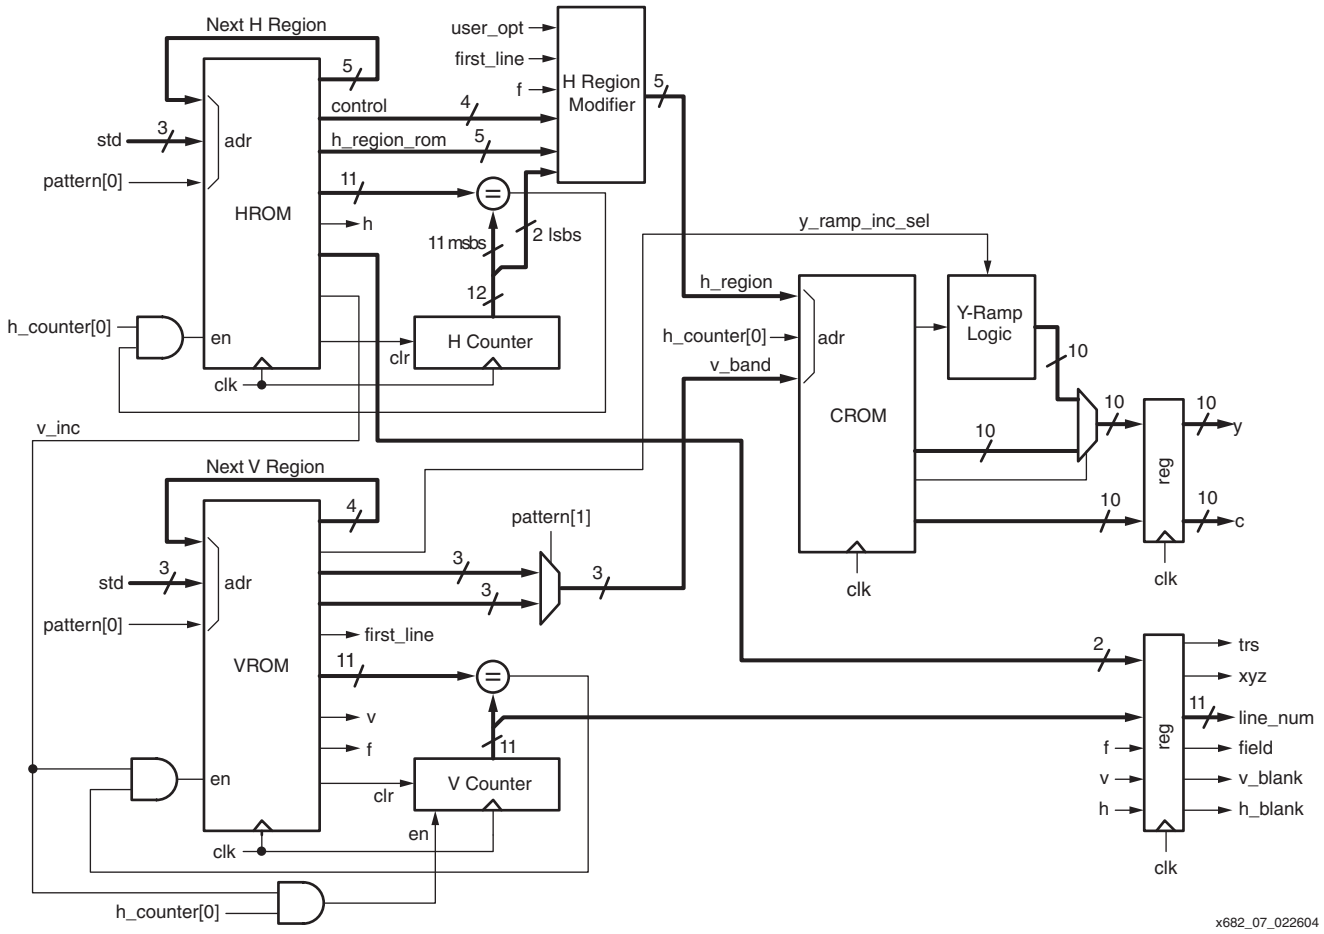

Figure 7: **Video Pattern Generator Block Diagram**

The video pattern generator consists of three main sections: the horizontal sequencer, the vertical sequencer, and the output generator. The horizontal sequencer keeps track of the current horizontal position and generates a horizontal region code indicating to the output section which horizontal region of the video pattern is currently being drawn. Likewise, the vertical sequencer keeps track of the current vertical position and generates a vertical band code indicating to the output section which vertical band of the video pattern is currently being drawn. The output section converts the horizontal and vertical coordinates of the video pattern into the luma and chroma video components appropriate for that portion of the pattern.

Horizontal Section

One block RAM (the HROM) is used as part of the horizontal sequencer. A horizontal counter increments every clock cycle, keeping track of the current horizontal sample count. The HROM outputs a next event value indicating the horizontal sample number when the sequencer should be advanced to the next region of the pattern. A transition from one horizontal region to another takes place whenever the video output value needs to change, such as when moving from one color bar to the next. Separate regions exist for the EAV, SAV, and horizontal blanking portions of the video pattern, as well. The LSB of the horizontal counter is not used in the comparison with the horizontal next event. A transition from one horizontal region to the next can only occur when the LSB of the horizontal counter is High.

The HROM implements a finite state machine (FSM). The current state register of the FSM is the block RAM's internal input register. The next region output of the HROM is the next state of the FSM. This value wraps back to the address inputs of the HROM and will be loaded into the current state register the next time the register is loaded. This only occurs when the horizontal counter matches the next event value from the HROM.

The HROM also outputs a value called h_region. This value indicates to the output section which region of the video pattern is currently active, horizontally. The h_region output from the HROM is modified in some cases to accommodate pattern modifications that the HROM can't handle, such as for the user option fields of RP 219-2002.

The HROM also outputs several other control signals. The h_clr output of the HROM causes the horizontal counter to clear to zero when the end of the line is reached. The v_inc output causes the vertical counter in the vertical section to increment once per line.

Vertical Section

The vertical section is very similar to the horizontal section. Central to the vertical section is another block RAM-based FSM called the VROM. A vertical counter keeps track of the current line number. The counter is compared to a vertical next event value from the VROM to determine when to advance the state of the FSM. As with the HROM, the VROM outputs a next state value that is fed back to the VROM address inputs. Whenever the vertical counter matches the next event value from the VROM, the next state value is loaded into the current state register in the VROM.

The VROM outputs a vertical band code that indicates which vertical "band" of the video pattern is currently active. Different bands are used for the different vertical sections of the pattern. For example, the RP 219-2002 pattern has four vertical bands: the 75% color bars on top, the two narrow rows in the middle, and the bottom mostly black and white band. Additional vertical regions are used to indicate the vertical blanking intervals.

The VROM actually outputs two vertical band codes, one for the 75% color bars and the other for all other video patterns. A MUX selects between these two codes based on the pattern selection inputs.

Output Section

The output section of the video pattern generator uses another block RAM (the CROM) to generate the luma and chroma values. The address inputs to the CROM are the vertical band (v_band) from the vertical section, the horizontal region (h_region) from the horizontal FSM, and the LSB of the horizontal counter. The vertical band and the horizontal region uniquely identify the current rectangular block of the video pattern. The CROM will output the luma and chroma components that correspond to this portion of the video pattern. Because 4:2:2 digital component video is being generated, consecutive samples of the chroma channel must alternate between containing a Cb' word or a Cr' word. The LSB of the horizontal counter tells the CROM which chroma word to output.

The CROM does not have the flexibility to generate the Y-Ramp portion of the RP 219-2002 color bar pattern so a Y-Ramp generator is added on the output of the CROM. Figure 8 shows the details of the Y-Ramp generator. The Y-Ramp generator is an accumulator with an output rounder. The accumulation register is initialized with an initial value called YRAMP_INIT when the CROM asserts its y_ramp_reload output. This occurs at the beginning of the Y-Ramp pattern. For each video sample after that, an increment value is added to the current value of the accumulation register and the sum is loaded back into the accumulator register. The increment value is different for the video formats with 1280 active samples versus those with 1920 active samples.

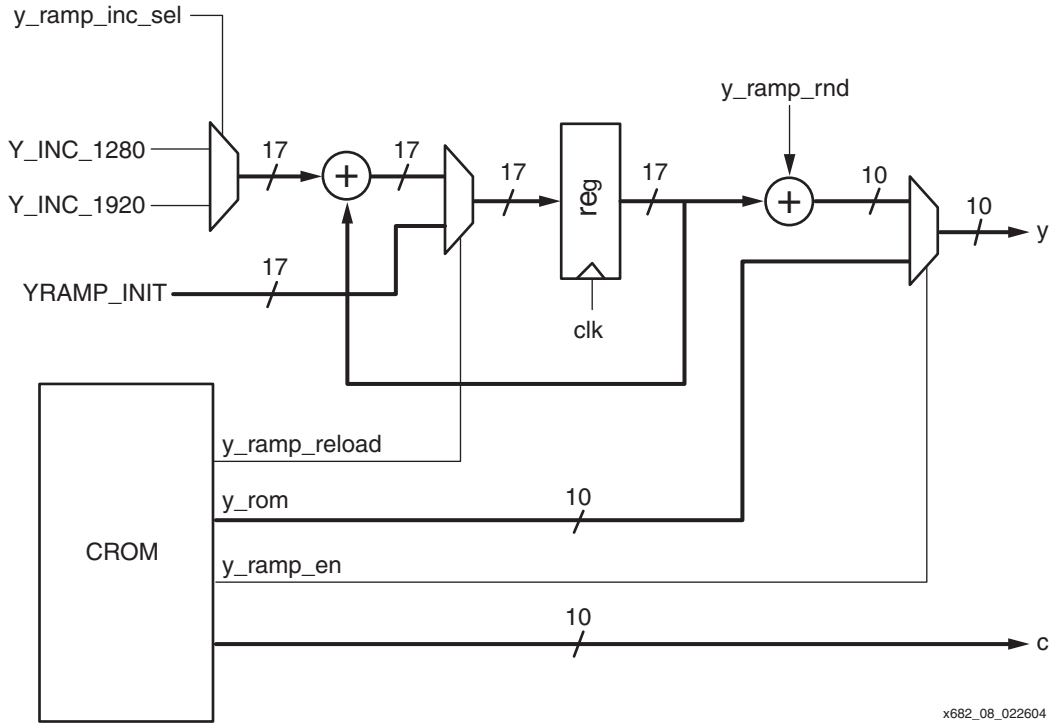

Figure 8: **Y-Ramp Generator**

To make the Y-Ramp pattern as linear as possible, the accumulation register and the increment value both have fractional bits. The output of the accumulation register is rounded to the nearest integer value before being output from the video pattern generator.

A MUX controls whether the luma output of the CROM or the rounded output of the Y-Ramp generator is used as the luma output. The MUX is controlled by the y_ramp_en output of the CROM so that the Y-Ramp generator output is only used during the Y-Ramp portion of the video pattern.

Other Outputs

The video pattern generator outputs several video timing signals:

trs: Asserted when the four-word timing reference sequence is output by the video pattern generator.

xyz: Asserted when the XYZ word of the timing reference sequence is output by the video pattern generator.

field: Indicates whether the current field is the first (field = Low) or second (field = High) field. This signal is always Low for progressive formats.

v_blank: Asserted during the vertical blanking interval.

h_blank: Asserted during the horizontal blanking interval.

The video pattern generator also outputs an 11-bit line_num value indicating the current video line number. HD-SDI requires that the line number be inserted into the two words that immediately follow EAV timing reference sequences. The video pattern generator does not insert the line number because this is only required when transmitting the video over HD-SDI. The line number is output from the video pattern generator so that logic in the HD-SDI transmitter can insert it into the video stream.

Initialization File Generation

The video pattern generator is a relatively simply design because all of the complexity of the video patterns is encoded into the block RAMs. The creation of these initialization files for the block RAMs is where all of the complexity occurs.

The initialization values for the three block RAMs are created by a Verilog file called `multigenHD_romgen.v`. This file, when compiled and run using a Verilog simulator, will create six different initialization files for each block RAM:

- Verilog simulation init file `(xxxx_sim.v)`
- Verilog XST synthesis file `(xxxx_xst.v)`
- Verilog Synplify synthesis file `(xxxx_syn.v)`
- VHDL simulation init file `(xxxx_sim.vhd)`
- VHDL XST synthesis file `(xxxx_xst.vhd)`
- VHDL Synplify synthesis file `(xxxx_syn.vhd)`

While there is no VHDL version of the `multigenHD_romgen.v` file, this Verilog file creates both Verilog and VHDL initialization files. These files can be inserted into the appropriate places in the video pattern generator design files.

The `multigenHD_romgen.v` file can produce initialization files for any subset of the nine different format groups listed in Table 2. The subset included in the initialization files can be easily changed and the mapping from the 3-bit standard input code to the selected video format can also be modified simply by changing the convert_std function at the end of the `multigenHD_romgen.v` file. This function consists of a simple case statement that maps the 3-bit standard code into a 4-bit video format code used by the ROM generation routines to create the initialization files.

To generate the initialization files, compile and load the multigenHD_romgen.v file into a Verilog simulator and then run the simulation to completion. The simulation will automatically stop when the files are completed. The `multigenHD_romgen` file has been tested with the current version of ModelSim.

Design Size

Table 3 lists the FPGA resources required for the multigenHD video pattern generator. These results were obtained using XST and ISE 6.1. This reference design has been simulated using ModelSim. It has also been tested in hardware using the Xilinx SDV demo board. [Ref 3]

Table 3: **Reference Design Results**

	LUTs	FFs	Block RAMS
multigenHD	119	100	3

Other Considerations

The block RAMs that are used as the three ROMs in this design are dual-ported. In the reference design, only one of the ports is used. However, it is possible to use the second port of each ROM to create a second, independent video pattern generator. The other logic in the video pattern generator must be duplicated for the second pattern generator, but using the same block RAMs for both pattern generators does save FPGA resources in applications that require two video pattern generators. The size of a dual video pattern generator can be estimated by simply doubling the LUT and flip-flop counts listed in Table 3.

The output values of the video pattern generator are not bandwidth limited—the rise and fall times are not shaped. When these patterns are displayed on an analog display, the sharp rise and fall times at the transitions between bars can cause ringing of the video image at these transitions. RP 219-2002 recommends that the rise and fall times be limited to the equivalent of 55 ns from 10% to 90% for all three components. If necessary, this can be done in the FPGA by filtering the three components with a digital filter. Note, however, that the RP 198-1998 checkfield should never be filtered. The values of this checkfield must enter the HD-SDI encoder exactly as they come from the video pattern generator.

Conclusions

This application note describes a reference design capable of generating common HDTV video test patterns for all 18 video formats supported by the HD-SDI standard. The reference design can produce the SMPTE RP 219-2002 color bar pattern along with two test patterns commonly used to test HD-SDI interface performance.

The reference design is based on three 18 kb block RAMs, allowing the video patterns and supported video formats to be changed by changing the block RAM initialization files. A Verilog file is used to create the block RAM initialization files. The use of block RAMs as the core of the video test pattern generator results in a small, efficient implementation.

References

1. Xilinx Application Note XAPP248, "Digital Video Test Pattern Generators" by John F. Snow
2. All the SMPTE standards referenced in this application note are available from The Society of Motion Picture and Television Engineers. These standards can be purchased at the SMPTE web site: http://www.smpte.org.
3. The Xilinx SDV Demo board is available from Cook Technologies (part number CTXIL103). Further information is available at http://www.cook-tech.com.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
03/06/04	1.0	Initial Xilinx release.