

This paper describes how Altera's 40-nm and 28-nm FPGAs are tailored to help deliver highly-integrated, HD studio equipment products. The paper provides an analysis of the performance requirements, resource utilization, and power consumption characteristics for the format conversion of multiple video channels. This is a common function for broadcast applications ranging from video capture cards to multiviewers, video walls, and A/V switchers. The paper also describes the architectural enhancements featured in Altera's 28-nm FPGAs that specifically improve their capability for broadcast applications.

Introduction

Increasing industry demand to deliver HD video channels requires studio equipment providers to deliver integrated products that provide the required bandwidth and processing power, while minimizing cost and power. Although some studio equipment providers resort to full custom ASICs, time-to-market pressure and development expense often rule out this option. Application-specific standard products (ASSPs) provide an alternative in some applications, but they can be inflexible and cannot provide high integration relative to shifting market demands.

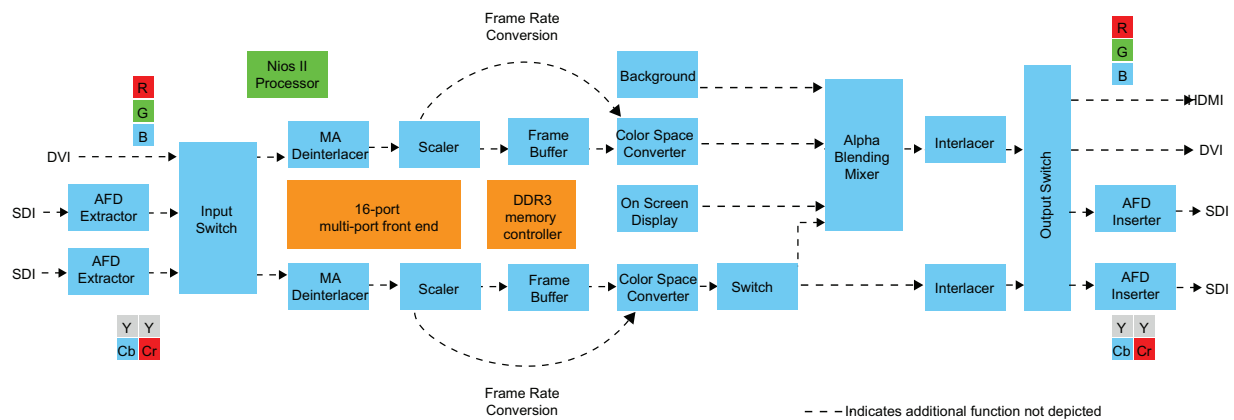
Against this backdrop, Altera offers its latest generation of 40-nm and 28-nm FPGAs tailored to deliver studio equipment developers higher integration and customization than ASSP-based systems, while avoiding the lengthy development times and costs of full custom ASICs.

Up/Down Cross Conversion (UDX) Requirements

The process of converting video prior to storage, encoding, or display can be described as up/down cross conversion (UDX). [Figure 1](#) shows a simplified block diagram of a 2-channel UDX design developed by Altera. This design has extensive functionality, in addition to simple format conversion, and correspondingly overestimates required gate resources for most applications. This design is used to analyze the fitness, performance, and power characteristics of Altera FPGAs implemented in studio equipment products.

The 2-channel UDX design ingests video over serial digital interface (SDI) or digital visual interface (DVI). This design can handle two SD-SDI, HD-SDI, or 3G-SDI progressive or interlaced input streams up to 1080p60, such as NTSC, PAL, 720p, 1080i, and 1080p. The Active Format Description (AFD) Extractor extracts code from the SDI channels to support dynamic clipping, scaling, and padding for bidirectional format conversion between 4:3 and 16:9 aspect ratios. Next, the input switch performs 4:2:2 to 4:4:4 chroma sampling conversion as required, which allows selection of two of the three input streams for input to the two video processing channels.

Figure 1. 2-channel Up/Down Cross Conversion (UDX) design developed by Altera



Within a video processing channel, a motion-adaptive (MA) deinterlacer deinterlaces the video input in 4:2:2 mode, double-buffering it in external RAM, one output frame for each input field. Following that, the video frames are scaled to the desired resolution and buffered in external memory for frame rate conversion. The converted image is then mixed with the second channel and logos before displaying the image over user-selectable output such as SDI, DVI, or HDMI.



The UDX design has been successfully implemented and demonstrated in hardware.

Calculating Resource and Memory Requirements

The memory bandwidth requirements for Altera's UDX design are determined by the deinterlacing stage and associated frame buffering. The per-channel device resource requirements for the UDX design are as shown in [Table 1](#):

Table 1. UDX Design Device Resource Requirements (Per Video Channel)

Resource	Minimum FPGA Resources	External RAM
Logic elements (LEs)	45K	N/A
Internal RAM (Mbits)	2.6	N/A
DSP (18x18 multipliers)	110	N/A
Transceiver channels	1 (SDI or DVI)	N/A
External RAM (Mbytes)	N/A	13.22

1080p Memory Bandwidth

The memory bandwidth requirements are defined by the maximum resolution video the channel must handle. Since the design handles resolutions up to 1080p Video, the following equation calculates the memory bandwidth required to buffer a 1080p video:

$$\begin{aligned} \langle \text{Each 1080p frame width} \rangle \times \langle \text{height} \rangle &= 1920 \times 1080 = 2073600 \text{ bits} \\ 2073600 \times 60 \text{ FPS} \times 2 \text{ color planes} \times 10 \text{ bit resolution} &= 2.48832 \text{ Gbps} \end{aligned}$$

Therefore, the minimum memory bandwidth required to write 1080p video is 2.48832 Gbps. However, the design must also account for the maximum size of word determined by the width of the memory interface. For the target FPGAs, a 64-bit memory interface is assumed, which yields a 256-bit word. To avoid splitting pixels, 12, 20-bit pixels per read or write are packed into a 256-bit word with 16 unused bits:

$$12 \text{ pixels} \times 20 \text{ bits} = 240 \text{ bits.}$$

Thus, the actual bandwidth required to read or write 1080p video without splitting pixels in a 64-bit memory interface can be expressed as follows:

$$2.48832 \text{ Gbps} \times (256/240) = 2.654208 \text{ Gbps}$$

Motion-Adaptive Deinterlacing Algorithm

The motion-adaptive deinterlacing algorithm requires one write at 1080i, plus either four reads at 1080i, or two reads at 1080p:

$$1 \text{ write @ } 1080i = 0.5 \times 2.654208 \text{ Gbps} = 1.327104 \text{ Gbps}$$

$$4 \text{ reads @ } 1080i \text{ or } 2 \text{ reads @ } 1080P = 2 \times 2.654208 \text{ Gbps} = 5.30816 \text{ Gbps}$$

$$\text{Total} = 6.635264 \text{ Gbps}$$

If the deinterlacer includes the motion bleed feature, the store and compare motion values of the current frame must be compared with stored values. The motion-adaptive deinterlacing algorithm also requires one write and one read of video motion values; the minimum bandwidth required for each read or write assuming 10-bit motion values is as follows:

$$1920 \times 1080 \times 60/2 \text{ FPS} \times 10 \text{ bits} = 0.622 \text{ Gbps}$$

At 10 bits per motion value, a total of 25 motion values can fit into a single 256-bit word. Taking into consideration the avoidance of splitting pixels across the 256-bit word, the bandwidth required becomes:

$$0.622 \text{ Gbps} \times (256/250) = 0.637 \text{ Gbps}$$

So, the memory bandwidth required for a single channel of motion-adaptive deinterlacing is:

$$6.635264 \text{ Gbps} + (2 \times 0.637 \text{ Gbps}) = 7.90953984 \text{ Gbps}$$

Similarly the bandwidth required for a framebuffer is calculated by adding memory requirements for writing and reading one 1080p frame:

$$2.48832 \text{ Gbps} \times (256/240) \times 2 = 5.308 \text{ Gbps}$$

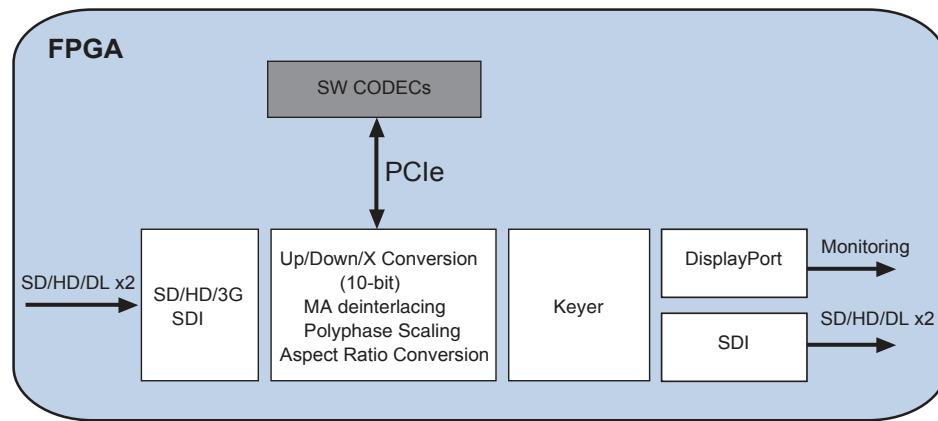
Hence the total memory bandwidth required per UDX channel equals the sum of memory bandwidth requirements of the deinterlacer and the frame buffer

$$7.90953984 \text{ Gbps} + 5.308 \text{ Gbps} = 13.21795584 \text{ Gbps, or } \sim 13.22 \text{ Gbps}$$

Implementing the UDX Design in 40-nm and 28-nm FPGAs

Consider a simple two-channel UDX design, common to capture cards, such as the one shown in [Figure 2](#).

Figure 2. PCIe Capture Card with Two-Channel UDX



The memory bandwidth requirements for the two-channel UDX design is calculated as follows:

$$2 \text{ channels} \times 13.22 \text{ Gbps} = 26.44 \text{ Gbps}$$

[Table 2](#) outlines the resources required for a 2-channel PCIe[®] capture card, including a DisplayPort output for monitoring, and a PCIe interface to transfer the video data to the host and access software codecs.

Table 2. FPGA Resources Required for 2-Channel PCIe Capture Card

Resource Type	Resources per Channel	Two Format Conversion Resources	DisplayPort and PCIe Interface Resources	Total Capture Card Resources
Logic element (LE)	45K	90K	12K	102K
Internal RAM (Mbits)	2.6	5.2	0.3	5.5
DSP (18x18 multipliers)	110	220	N/A	220
Transceiver channels	2 (SDI or DVI)	4 (2 input, 2 output)	4 (DisplayPort) plus 4 (PCIe Gen2x4) or 8 (PCIe Gen1x8)	12 or 16

[Table 3](#) below shows the target 40-nm and 28-nm FPGAs that are the best fit for the capture card design, as well as the relevant device resource counts. For the maximum memory bandwidth, symmetric interfaces (that is, at least two interfaces of same width and speed) are noted because sometimes the FPGAs can support higher memory bandwidth with additional interfaces of different data widths, and/or speeds. However, since this situation is often not desirable or practical, only the maximum bandwidth with symmetric interfaces is shown. Both FPGA options easily meet the memory bandwidth requirement of 26.44 Gbps, as indicated by [Table 3](#).


Table 3 also indicates the nature of memory interface support for the specified target devices. Altera's 40-nm FPGAs offer external memory interfaces via soft memory controllers, implemented in the user-programmable logic and memory portions of the device. These soft controllers have been demonstrated and tested with the UDX design in actual hardware, and they have proven to deliver the required efficiency and resulting bandwidth required.

In the 28-nm Arria® V FPGA, the memory interface is implemented in a hard memory controller. This hard memory controller is based on the proven soft memory controller, and is designed to provide even higher efficiency, along with easy, built-in timing closure.

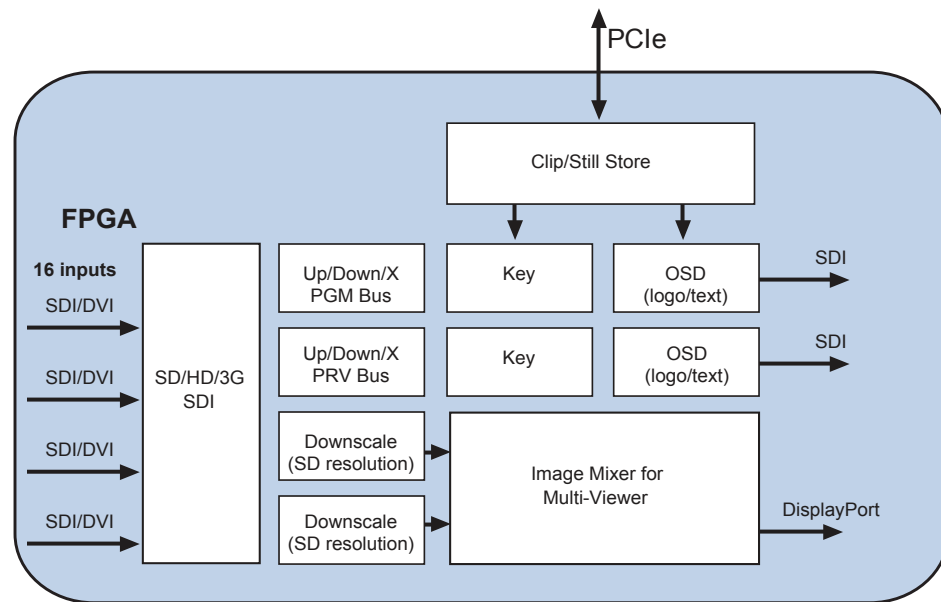
Table 3. FPGA Resources and Total Power Consumption

FPGA Resource	Arria II GX (40nm)	Arria V (28nm)
Target device	2AGX190	5AGXA3
Logic elements (LEs)	190K	150K
Total Memory (Mbits)	9.9	10.4
Max 18x18 multipliers	656	792
Max transceiver channels	16	12
Max memory bandwidth with symmetric interfaces	51.2 Gbps (soft controller)	136.4 Gbps (hard controller)
PCIe hard IP support	Up to Gen1x8	Up to Gen2x4
Capture card total power consumption	10.8 watts	5.8 watts

The last row in **Table 3** indicates the total power consumption for the capture card design as implemented in each device. This power is calculated using the PowerPlay Early Power Estimator (EPE) tool. Both FPGA options provide the lowest total power at their respective process nodes, delivering significant benefits for the increasingly power-sensitive end markets in the broadcast space.

 For more information about the EPE tool, visit the [PowerPlay Early Power Estimators \(EPE\) and Power Analyze](#) website.

A larger design based on the UDX design can better demonstrate the full integration capabilities of the most advanced FPGAs. For example, a 16-input, 8-channel A/V switcher, as shown in **Figure 3**.

Figure 3. 16-input AV Switcher with 8-Channel UDX

The design shown in [Figure 3](#) requires only a single advanced FPGA to implement. However, this design would require multiple ASSPs, along with the associated additional board space, power consumption, and higher design complexity.

The first step in implementing this design in a single FPGA is to calculate the memory bandwidth required for the 8 channels of UDX as follows:

$$8 \text{ channels} \times 13.22 \text{ Gbps} = 105.76 \text{ Gbps}$$

[Table 4](#) below outlines the resources required for a 16-input, 8-channel switcher, including a DisplayPort output for monitoring, and a PCIe interface to transfer the video data to the host and obtain clips and still images.

Table 4. Required FPGA Resources for 16-input, 8-Channel A/V Switcher

FPGA Resource	Resources per Channel	Resources per 8 Channels	DisplayPort and PCIe interface	16-Input, 8 Channel AV Switcher Resources
Logic elements (LEs)	45K	360K	12K	373K
Internal RAM (Mbits)	2.6	20.8	0.3	21.1
DSP (18x18 multipliers)	110	880	N/A	880
Transceiver channels	2 (SDI or DVI)	24 (16 input, 8 output)	4 DisplayPort plus 8 PCIe (2 Gen2x4, Gen2x8)	36

Table 5 shows the target 40-nm and 28-nm FPGAs that are the best fit for the 16-input, 8-channel A/V switcher design, as well as their relevant device resource counts. As described, only symmetric interfaces are used to determine the maximum memory bandwidth, and both options easily meet the memory bandwidth requirement of 105.76 Gbps.

Table 5. FPGA Device Resources and Total Power Consumption for 16-Input, 8-Channel A/V Switcher

FPGA Resource	Stratix IV GX (40nm)	Arria V (28nm)
Target device	EP4SGX530	5AGXB7
Logic elements (LEs)	531.2K	500K
Total memory (Mbits)	27.3	23.7
Max 18x18 multipliers	1040	2278
Max transceiver channels	48	36
Max memory bandwidth with symmetric interfaces	136.4 Gbps (soft controller)	136.4 Gbps (hard controller)
PCIe hard IP support	Up to Gen2x8	Up to Gen2x4
A/V Switcher total power consumption	22.4 watts	15 watts

In addition to implementing this complex design in a single chip, the FPGA options deliver the lowest total power of any FPGA implementation at their respective process node, thus providing the most attractive solution at every product generation. In addition, designers benefit from an easy migration path to next generation FPGAs, since the underlying technology of the UDX design and associated memory controller architecture is consistent across FPGA generations.

28-nm FPGA Optimizations for Broadcast Applications

In addition to providing consistency at the algorithm and implementation level, Altera also made specific architectural enhancements in its 28-nm FPGAs to better meet the needs of broadcast applications.

Optimized Video Embedded Memory Blocks

Altera configured its embedded memory blocks to efficiently and precisely accommodate 10-bit video data. Accordingly, Altera offers embedded memory blocks in its 28-nm devices that can be configured with widths in increments of 10 (that is, x10, x20, and x40) without wasting bits. Altera's broadcast-focused optimization contrasts with older FPGA architectures in which the embedded memory blocks are arranged in 18- and 36-bit widths, which results in inefficiencies, wasted memory, and the use of larger devices to obtain the required memory resources.


Variable-Precision DSP Blocks

Another broadcast-focused optimization is the introduction of variable-precision DSP blocks. These blocks can implement multipliers of various precisions, including 9x9, 18x18, and 27x27. In addition, designers can cascade the variable-precision DSP blocks to efficiently implement higher precision multipliers. For example, the UDX design requires multiplications of up to 10x16 (10 bits x up to 16-bit coefficients). Each variable-precision DSP block can implement two multipliers of 18x18 precision, which covers the 10x16 maximum precision required by the UDX design. In older FPGA architectures, a 10x16 multiplication may require a full DSP block, and older DSP blocks cannot be decomposed into lower precisions, which results in inefficient implementation utilization of more FPGA resources than necessary.

Lowest Power Transceivers

Another important optimization is the reduction of transceiver power. Many broadcast applications require increasingly more SDI channels, and therefore more transceiver channels. The benefits of higher integration are severely mitigated if the resulting design consumes high amounts of power that requires additional cooling costs, or produces a less competitive product. Altera is continuing its trend of transceiver power reduction by reducing the power-per-channel of its transceivers at the 28-nm node. This reduction allows designers to integrate more transceiver channels into a single device, while maintaining or reducing their thermal budget. [Figure 4](#) shows the historical trend of power-per-transceiver across three generations of FPGAs, and demonstrates Altera's commitment and ability to reduce transceiver power. This commitment reflects a decade of internal transceiver expertise that is unmatched in the industry. The significant reduction in transceiver power contributes to Altera's ability to provide the lowest total power FPGAs.

Figure 4. Historical Trend of Transceiver Power-Per-Channel in FPGAs



Conclusion

The bandwidth and power challenges faced by broadcast-equipment developers can be met with today's FPGAs. Equipment developers leveraging FPGAs can benefit from highly-integrated hardware-accelerated video processing and vendor-provided IP frameworks. These frameworks provide common video building blocks while enabling designers to focus on proprietary functions. The most comprehensive FPGA offerings combine low-power approaches and proven video processing techniques to minimize risk, while providing a clear roadmap to even more advanced FPGAs with broadcast-specific architecture enhancements and optimizations for even lower power.

Further Information

- *Meeting the Low Power Imperative at 28nm*
<http://www.altera.com/literature/wp/wp-01158-low-power-28nm.pdf>
- *Reducing Power Consumption and Increasing Bandwidth on 28-nm FPGAs*
<http://www.altera.com/literature/wp/wp-01148-stxv-power-consumption.pdf>

Acknowledgements

- Girish Malipeddi, Senior Technical Marketing Manager, Altera Corporation.
- Martin S. Won, Senior Member of Technical Staff, Altera Corporation.

