

Satisfying the Demand for Rapid Feature Enhancement in Consumer Display Products

Introduction

Developers of consumer display products (i.e., high-definition televisions (HDTVs), monitors, and projectors) face a daunting design challenge. The application-specific standard products (ASSPs) that have been used in these products historically are built with custom ASIC development cycles, typically three to five years. However, the consumer market demands new products and features on a yearly basis. The result is that consumer product developers cannot rely solely on fixed-function ASSPs to deliver the frequency of innovation demanded by their customers.

Increasingly, these developers are turning to FPGAs to augment standard chipsets, thereby adding advanced features and capabilities in a shorter time frame (usually six to nine months). Image enhancement features, like frame rate conversion, upscaling, and noise reduction, are being implemented in FPGAs and applied in consumer display applications.

This paper describes how designers can take advantage of this approach with Altera® Cyclone® III FPGAs, which meet the performance, I/O, signal processing, memory, and cost efficiency requirements of this application space.

Proliferation of HDTVs Relative to Available HD Content

The numbers of HDTVs are increasing rapidly, and for the coming years, analyst projections foresee double-digit growth on a percentage basis. Accordingly, the number of HD broadcast networks around the world is increasing the amount of HD content created. Despite this, the majority of material viewed on televisions is still far below HD resolution, including standard-definition (SD) cable, broadcast, and satellite transmissions, as well as content from storage media such as DVDs.

This is becoming a major problem for the broadcast industry. This SD content must be de-interlaced and upscaled from 480 lines to 1080 lines without introducing fringing, noise, or artifacts. Broadcasters are making substantial investments in equipment to convert SD content to content suitable for HD transmission. Also, with the worldwide rollout of HD services, there is an increasing need to convert to and from the 60-Hz refresh rate used in much of the Americas, Japan, and Korea, as well as the 50-Hz refresh rate used in the rest of the world. This was hard enough at SD, but HD is much less forgiving and any artifacts show up more clearly on big screens. Converting 60-Hz SD to 50-Hz HD is even more difficult, and again broadcasters must make large expenditures for frame-rate converters to tackle this challenge.

On the consumer side, manufacturers of consumer displays and home theater peripherals have the opportunity to leverage upscaling and frame-rate conversion technology to deliver product-differentiating image quality. Features like this can fulfill the potential of high-resolution HDTVs by improving the viewing experience for the majority of available content. However, the improvements must be significant and delivered at the lowest possible cost in order to gain traction with consumers.

The problem is that, over the last ten years, the development of the algorithms behind upscaling and frame rate conversion has not kept up with that of HDTV. Bilinear and bicubic interpolation and pixel filtering struggle with the move to HD, and the next-generation motion-adaptive interpolation requires significant processing power and large quantities of expensive custom chips. [Figure 1](#) shows how current motion-adaptive upscaling algorithms based on motion-adaptive algorithms can produce jagged edges, as shown on the left, compared to more advanced algorithms (on the right) that focus on the full 3D directional analysis of the video, producing a better quality of results.

Figure 1. Results of Current Upscaling Techniques (left) vs. More Advanced Upscaling Techniques (right)



Measuring and predicting spatial motion can only go so far. One problem that cannot be solved in this manner is when two occluding blocks are moving in different directions. In this case, calculating the average motion creates an artifact. Alternate schemes to blur the results are not ideal and take additional processing power to detect and implement.

Tremendous efforts are going into solving this problem, partly to improve the quality of the content—particularly for graphics and video overlays, which motion compensation cannot handle very well—but also to handle text. These techniques leave moving text on the screen looking blurred or jumping around on HD screens. This effect is shown in Figure 2, where a poor image created by traditional upscaling algorithms appears on the left, and the same image upscaled using a more advanced approach appears on the right.


Figure 2. Visual Quality of Moving Text Produced by Traditional Upscaling Algorithms (left) vs. Advanced Upscaling Algorithms (right)



Innovative Approaches to Scaling and Frame Rate Conversion

Traditional motion-adaptive algorithms compute a missing pixel with a spatial or time interpolation from the previous and next pixels of the same location. Let It Wave has developed a new approach based on bandlets. Instead of concentrating on the motion within the image, the bandlet approach looks in both space and back over previous frames for the best pixels without flicker (vibrations in time) or jaggies (vibrations in space). By searching over a full space-time neighborhood for a directional interpolation, the bandlet approach minimizes the resulting total variation, as shown in Figure 3.

Figure 3. Motion-Adaptive Approach to Upscaling (left) Differs From the Bandlet Approach (right)



The bandlet approach is completely different from the way upscaling currently takes place, providing results that are better than the high-end, leading-edge systems available today. The bandlet approach therefore lends itself well to implementation in a relatively small, and therefore cost-effective, chip, and supports the whole range of upscaling up to 1080p in progressive displays of any size. Suitable for frame rate conversion of between 50 Hz and 60 Hz for broadcast equipment, and up to 120 Hz for large panel displays, the bandlet approach's small size makes the technology viable for high-volume consumer applications.

Implementing Video/Image Enhancement Algorithms for Consumer Products

Developers of consumer display products face an opportunity to leverage advances in image enhancement technology, like Let It Wave's bandlet approach, that offer differentiating features, while using the right implementation vehicle to minimize costs for the price-sensitive consumer market.

In many consumer products, video and image enhancement functions like upscaling are often performed by fixed-function ASSPs. Unfortunately, as described earlier, ASSP developers are challenged to meet the requirements of the consumer market, which demands new features and capabilities more rapidly than their custom ASIC-based product development cycles allow. It is also difficult for consumer product manufacturers to differentiate offerings if they rely solely on fixed-function components that are available to the entire developer community.

For some consumer applications and many other end applications, custom video and image enhancement functions are traditionally implemented in digital signal processors. This is because these functions are arithmetic-intensive, and there is a large amount of support for and experience in implementing these functions in digital signal processors. For HDTV applications, however, the processing requirements can be quite high, so the cost of a digital signal processor that delivers sufficient processing power can be prohibitive for cost-sensitive consumer products.

The core of the difficulty with digital signal processors in these applications is twofold. First, digital signal processors are essentially serial machines, processing one element of the signal chain at a time. In the case of some high-end digital signal processors, a small number of instructions can be processed simultaneously giving a small degree of parallelization. However, the cost can be many times that of a non-parallel version. In other applications, multiple digital signal processors are utilized to obtain true parallel processing, but the costs for hardware, and therefore power consumption, are quick to rise accordingly.

These issues have led designers to use Altera FPGAs to take advantage of the ability to convert a cascaded set of operations into a parallel structure that operates at much higher throughput than is possible with digital signal processors of equivalent cost. This is a central advantage of the programmable logic technology—the ability to speed up an algorithm by making the process truly parallel.

In a typical digital signal processing (DSP) application, engineers have few options to increase performance beyond writing pipelined assembly language (if the digital signal processor supports it), or replacing the digital signal processor with a higher frequency model. With the Altera approach, the hardware, as well as the software, can be

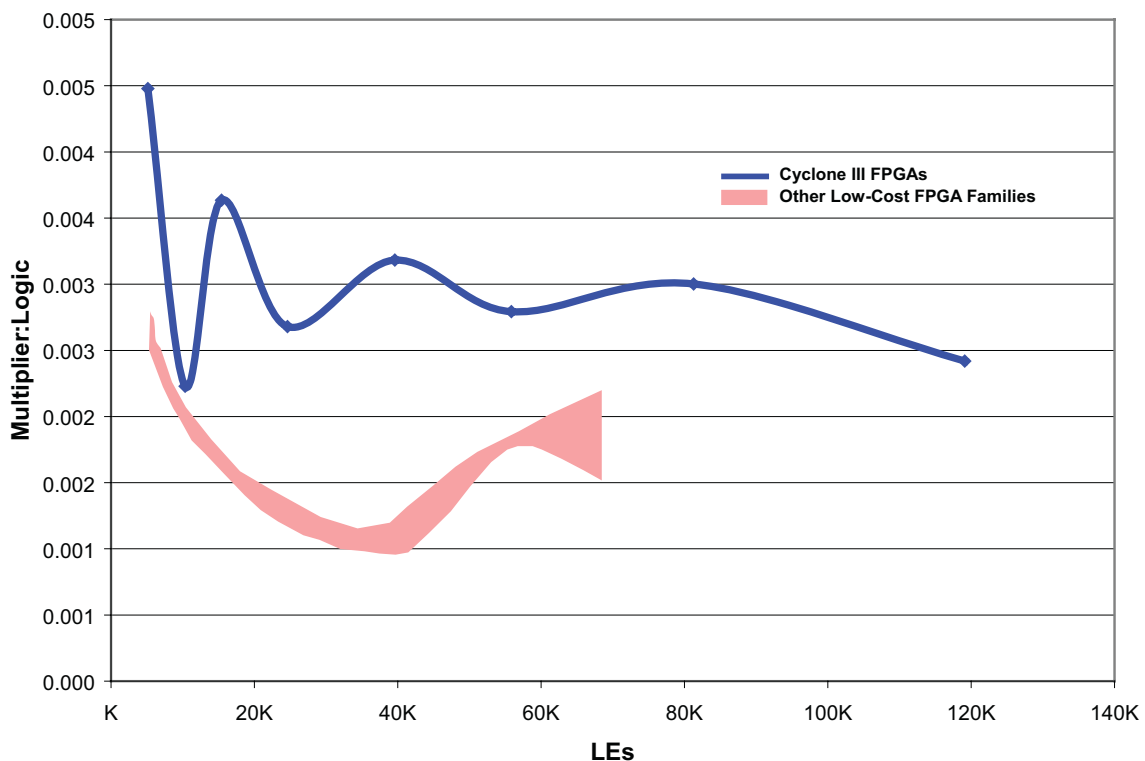
simultaneously optimized. Now the designer has a three-dimensional optimization space available: code, higher performance device, and hardware optimization.

With the nearly exponential price-versus-performance curves for many hardware processors, designers were left with usually only one option: code optimization. As anyone who has had to optimize code knows, that avenue is quickly exhausted. The broad flexibility provided by Altera's programmable solutions can help create systems that were previously impossible to design either because of time and cost of goods, or because traditional digital signal processors could not handle the calculation load. With Altera FPGAs, the engineer has one more degree of freedom in the design process—hardware acceleration—with which the developer can often achieve the best price/performance combination for any custom DSP operation.

In addition to ease of parallel processing, modern FPGAs offer embedded multipliers to implement DSP functions efficiently. Without a generous amount of embedded multipliers, arithmetic functions such as scalars and filters consume a high proportion of general-purpose logic relative to the other resources in the device, resulting in inefficient device utilization and potentially higher cost. As a result, FPGAs that offer a high multiplier-to-logic ratio are better suited for image and video enhancement functions.

To provide the best support for image and video enhancement functions, Altera Cyclone III FPGAs offer the highest multiplier-to-logic ratio, as shown by Figure 4, which shows that Cyclone III FPGAs offer the highest multiplier-to-logic ratio across the full range of logic densities, compared to other low-cost programmable logic families.

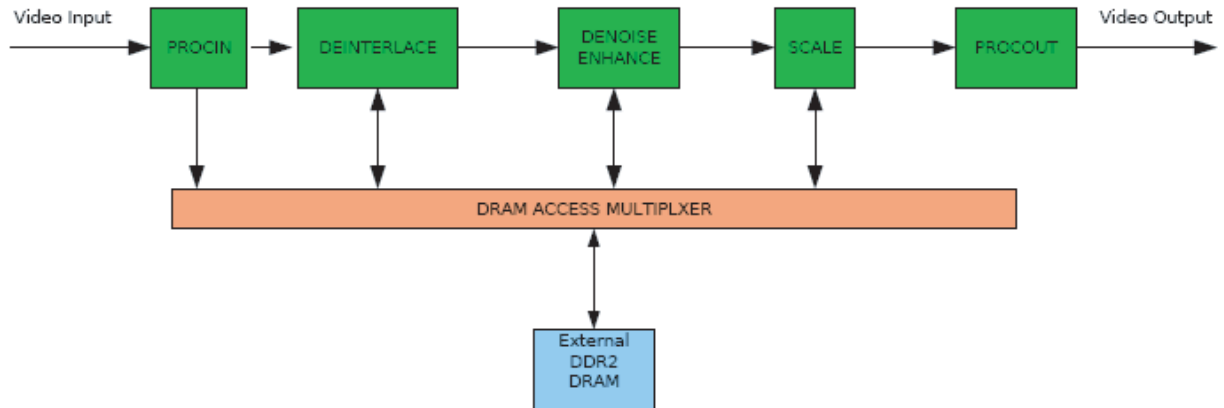
Figure 4. High Multipliers-to-Logic Ratio in Cyclone III Devices Compared to Other Low-Cost FPGAs



In addition to a high number of embedded multipliers, the ideal FPGA architecture for these functions requires large amounts of embedded memory. Image and video processing requires large amounts of memory bandwidth, with over 200 Mbytes per second needed for 1080p HDTV processing. The benefits of embedded memory can be illustrated by

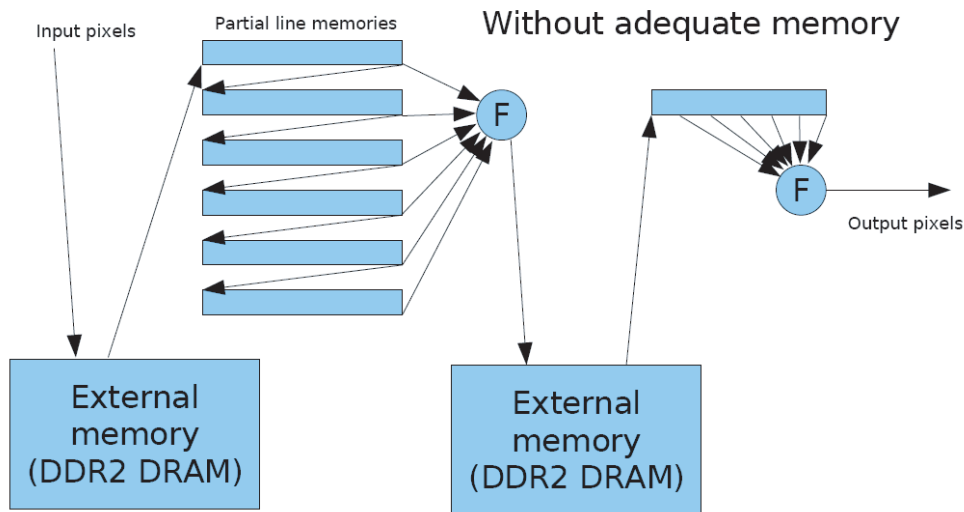
an example using Let It Wave's technology: [Figure 5](#) shows a block diagram of how video data flows through a system employing Let It Wave's image enhancement algorithms for de-interlacing, noise reduction, and scaling. The blocks in green as well as the DRAM access multiplexer are implemented in an FPGA.

Figure 5. Data Flow Through Let It Wave-Based Subsystem for De-Interlacing, Noise Removal, and Scaling




For the scaling stage, input pixels arrive in raster order, and without adequate internal memory, the pixels must be reordered through external memory and reread as partial lines before presentation to the vertical scaler. The pixels must also be reordered a second time to restore raster order before being sent to the horizontal scaler (see [Figure 6](#)).

Figure 6. Without Adequate Embedded Memory Within the FPGA, Video Enhancement Functions Such as Advanced Scaling Can Require Several Passes Through External Memory



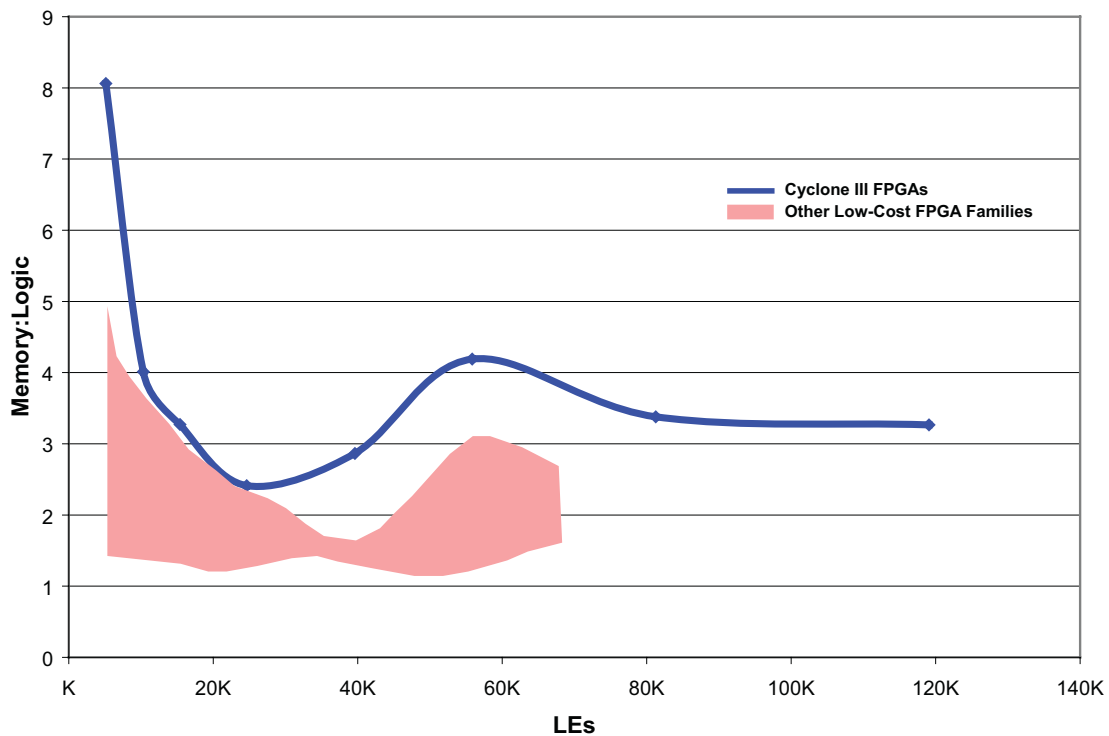
FPGAs that offer adequate embedded memory can buffer full lines of video, avoiding the need to reorder pixels through external memories. This results in lower latency and a potentially reduced implementation cost through the elimination of external memory devices. [Figure 7](#) shows how the pixels flow through Let It Wave's scaler—implemented in Cyclone III FPGAs, which offer sufficient embedded memory to eliminate the path through an external memory device. Implementing the Let It Wave scaler in Cyclone III FPGAs reduces the external memory requirements from four DDR2 DRAM devices to three, when compared with other FPGA architectures with less embedded memory.

Figure 7. With Adequate Embedded Memory Within the FPGA, External Memory Requirements Can Be Reduced or Eliminated



To deliver the best support for image and video enhancement functions, Altera Cyclone III FPGAs offer the highest memory-to-logic ratio across the full range of logic densities, compared to other low-cost programmable logic families, as shown by Figure 8.

Figure 8. High Memory-to-Logic Ratio in Cyclone III Devices Compared to Other Low-Cost FPGAs

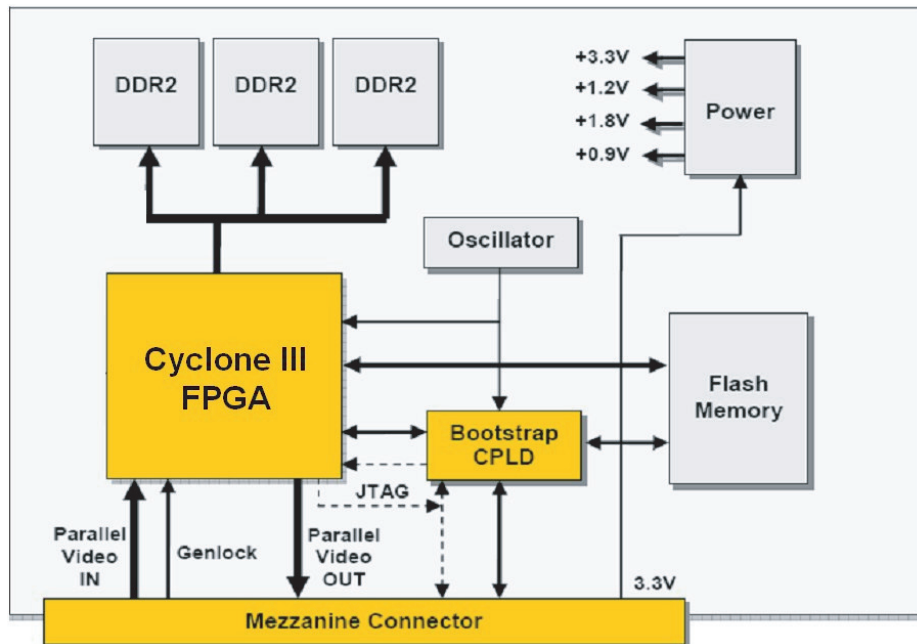


The Most Cost-Effective Delivery of Innovative Image/Video Enhancement

For the consumer market, Cyclone III FPGAs offer a memory- and multiplier-intensive architecture that is well suited to implementing image enhancement algorithms. In addition, as the world's first low-cost FPGAs developed on a 65-nm process, Cyclone III devices deliver the cost structure that opens up the market for integrating this technology into televisions and LCD panels. Consumer display makers already use Cyclone FPGAs for timing control, and there is the opportunity to integrate this alongside the bandlet upscaling and frame-rate conversion. This would allow panel makers to deliver the 120-Hz refresh rates that are the next step for HD consumer displays.

Figure 9 shows the diagram of a module based on a Cyclone III device that employs Let It Wave's technology for video enhancement. The module is designed to be easily integrated into customers' end products with the use of a mezzanine connector, and includes sufficient external DDR2 DRAM for HDTV operations, flash memory for holding FPGA configuration data, a CPLD for configuration management, an on-board oscillator for clock generation, and a voltage regulator for providing all of the needed power rails.

Figure 9. Cyclone III FPGA-Based Module for Implementing Let It Wave's Bandlet-Based Video Enhancement Algorithms



A variant of the technology is also available for the “prosumer” market, for upscaling home theater projectors and DVD players that need additional on-screen display (OSD) capability, which can be added to the FPGA easily. For any of these products, an FPGA-based design allows the developer to provide a range of products with a single design because FPGA reconfigurability allows the opportunity to add custom elements for without increasing the cost of the hardware.

Conclusion

The combination of innovative and disruptive image and video enhancement approaches, like Let It Wave's bandlet method with the low-cost Cyclone III FPGA family optimized for delivering these functions, presents several opportunities for developers of consumer display products and their suppliers. Flat panel and home theater projector manufacturers can deliver product-differentiating features cost-effectively and in the shortest amount of time. ASSP developers can likewise satisfy this demand by supplementing their chipsets with Cyclone III FPGA-based companion devices to provide frame rate conversion and the highest quality upscaling. This approach enables ASSP developers to preserve their initial investment while rapidly offering new capabilities within a development period that is impossible to achieve with custom ASICs.

Acknowledgements

- Stephane Mallat, President, Let it Wave
- Herve Mer, European Segment Marketing Manager for Broadcast and Consumer, Altera Corporation
- Martin S. Won, Senior Member of Technical Staff, Altera Corporation



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
<http://www.altera.com>

Copyright © 2007 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.