

AN FPGA-BASED JPEG 2000 DEMONSTRATION BOARD

Tom L. Woolston, Niel Holt, Gail Bingham, Glen Wada
 Space Dynamics Laboratory/Utah State University Research Foundation
 1695 North Research Parkway, North Logan, Utah 84341; (435) 797-4600
tom.woolston@sdl.usu.edu, niel.holt@sdl.usu.edu, gail.bingham@sdl.usu.edu, glen.wada@sdl.usu.edu

ABSTRACT: The Space Dynamics Laboratory has developed a hardware-based JPEG 2000 image compression solution and packaged it in a demonstration board. The board implements both Tier1 and Tier2 JPEG 2000 encoding in two Xilinx Virtex II FPGAs. The FPGA design was built as a first step toward developing JPEG 2000 image compression hardware that could be used for remote sensing on the ground, in the air, or in Earth orbit. This board has been used to demonstrate the power and flexibility of the JPEG 2000 standard in hardware, compressing both 8-bit and 12-bit grayscale images based on decoded image quality as well as output bit rate control. Images have also been compressed in both lossless and lossy modes. The board produces a JPEG 2000 file that includes all header and packet information needed to decode the output file. This paper will present the electrical design of the board and data flow. An application board of reduced size that maintains the same level of compression functionality has been conceptually designed.

INTRODUCTION

JPEG 2000^{1,2,3} is a relatively new image compression standard from the Joint Photographic Experts Group (JPEG). This standard is finding more application in the communication world as it offers many advantages over its predecessor, such as compressed file flexibility, improved image quality for a given compression ratio, good bit rate control, and ease of use. All of these features were the grounds for the development of a demonstration hardware compression system that uses the JPEG 2000 standard.

The board that was developed in this effort was not designed to fly in air or space, but rather was developed to understand JPEG 2000 image compression technology from a hardware perspective. Initial development issues included questions about what kind of effort would be needed to run the JPEG 2000 algorithm in a hardware environment. Throughput, firmware size, and understanding which features could be easily developed and which ones would require more effort were also addressed early in the project. Board size and power utilization were also issues to be addressed.

The board was designed with a peripheral component interconnect (PCI) bus interface for straightforward communications with a host computer. However, it should be noted that the board's two field programmable gate arrays (FPGAs), which contain the algorithms, are not in any way dependant on PCI-host

interfacing. These FPGA designs can be used with other communication busses, either standard or custom, depending on the application.

The major design objectives for the SDL JPEG 2000 demonstration board were to:

1. Generate a complete JPEG 2000 byte stream in hardware that could be immediately decoded with a standard decoder or further manipulated in software to enhance certain image features.
2. Use the Tier2 rate distortion algorithm for lossy compression based on bit rate control or output image quality.
3. Design an FPGA-based hardware solution, implementing as much JPEG 2000 functionality as possible in the FPGAs.
4. Provide a test bench for studying JPEG 2000 performance in a hardware environment.

Other considerations, such as board size were not primary goals for this design; rather, the board was designed with the intent that application boards would be much smaller and tailored to meet specific requirements. Also, the immediate applications for this hardware did not allow long-term storage of the compressed file in memory.

This paper presents the electrical design of the JPEG

2000 hardware, including data flow and class of parts. Algorithms used in this design and relevant experimental results are explained in another resource.⁴

SYSTEM FEATURES

Specifications for the JPEG 2000 hardware solution are given in the following table.

Image Size	256 x 256 pixels up to 2048 x 1024 pixels.
Pixel Depth	8-bit or 12-bit gray scale
Tile Size	128 x 128 pixels
Code Block Size	32 x 32 pixels
Output Quality Layers	1
Wavelet Types	5/3 and 9/7
Maximum Output File Size	512 Kbytes
Throughput	Bench mark: 10 Mps with 2048 x 1024 8-bit pixels
File Header Type	Unwrapped J2K or JP2
Output File Type	Sequential tile or tile part
Compression	Lossless or lossy (based on bit rate control or image quality)

JPEG 2000 DEMONSTRATION BOARD

Figures 1 and 2 show the front and back sides of the JPEG 2000 demonstration board. This board is approximately 4.2 in x 12.2 in. in size, and as explained in the introduction, was not designed for use in any particular system, rather just to develop the code for the FPGAs. Application boards can be designed much smaller for use in image compression systems such as those used in aircraft or spacecraft environments.

The JPEG 2000 demonstration board is designed with a PCI interface, a simple method for communication with a host computer.

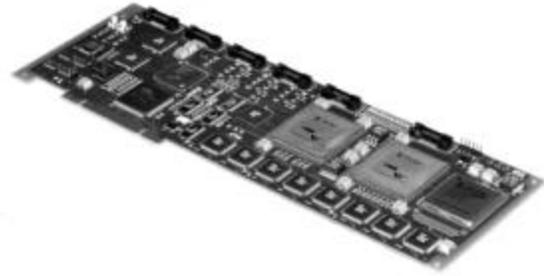


Figure 1. JPEG 2000 Board Front

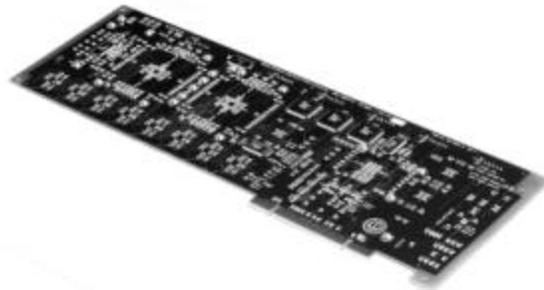


Figure 2. PCI Board Back

HARDWARE SYSTEM

Hardware Overview

From a system point of view, two Xilinx Virtex II FPGAs handle the main processing. Supporting the FPGAs are dual-port random access and first-in first-out (FIFO) memories. Figure 3 illustrates the interconnectivity of the processing and memory elements. The input memory holds lines of the input image and is used to convert these lines to input tiles. The FIFOs between the FPGAs hold code block bytes and portions of the JPEG 2000 output byte stream. The output memory holds the output byte stream while it is waiting to be read by the host computer. The PCI interface FPGA is used to format data that is transferred to and from the PCI driver chip.

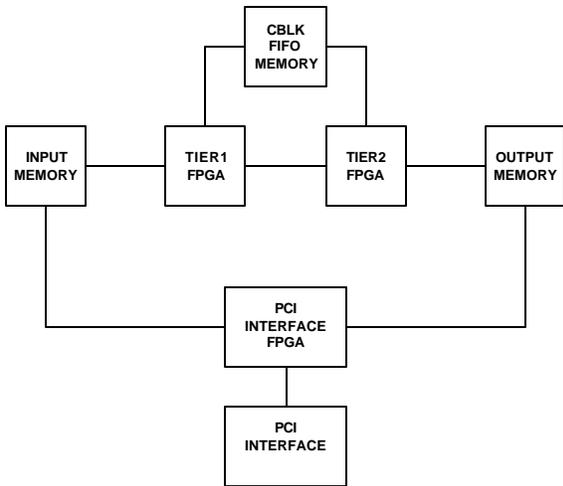


Figure 3. JPEG 2000 PCI Card Block Diagram

Input Memory

The input memory is made up of one dual-port RAM and eight FIFOs as shown in Figure 4. Horizontal image lines are loaded into the dual-port RAM from the host computer over the PCI bus. A horizontal image line is then divided into tile lines by moving the first 128 pixels of the first image line into a FIFO. The next 128 pixels are then placed into the next FIFO, and so on. This process continues until the last pixels in the line are placed into a FIFO (depending on the horizontal length of the image, these last pixels may or may not be placed into the last FIFO). Figure 5 shows an example of this process. Here, a horizontal image line of 480 pixels is divided into tile lines of 128 pixels with the exception of the last tile line. This process of dividing up the image lines into tile lines continues until all the image lines are loaded into the FIFOs. At this point, the image is in tile format.

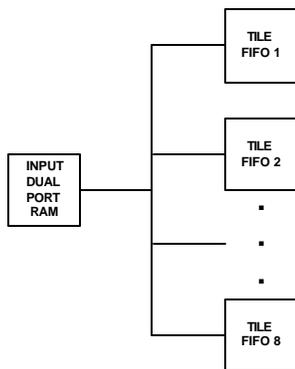


Figure 4. Input Memory Block Diagram

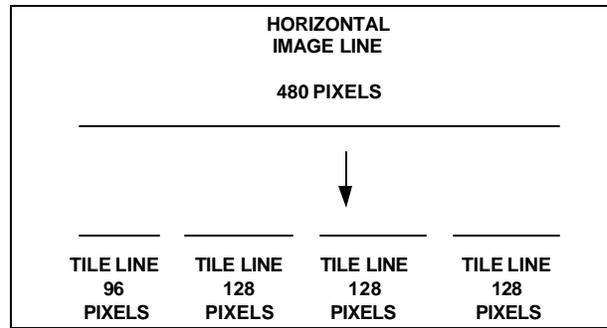


Figure 5. Lines-to-Tiles Scheme

For the compression process, a single tile is read from each FIFO. In the case of a 128 x 128 pixel tile, 128 tile lines are read from the first FIFO. The second tile is read from the next FIFO and so on, until a row of tiles is processed. The next row of tiles is then read and the process continues until all of the tiles are read.

Tier1 FPGA

The first compression process is performed in the Tier1 FPGA. The Tier1 FPGA compresses one tile at a time through a process that includes a two dimensional wavelet transform, quantization, and arithmetic encoding. The arithmetic encoding produces compressed code block bytes as well as code block attributes (these will be used in the Tier2 compression). Following the encoding, the data is reformatted for transfer to the code block FIFOs and the Tier2 FPGA.

Tier2 FPGA

The second compression process, a lossy compression, occurs in the Tier2 FPGA and uses the JPEG 2000 rate distortion compression algorithm. This algorithm can optimize the compression process by taking the content of the tile as well as the requested compression ratio into account. The algorithm can accommodate the optimization of one tile or a group of tiles (up to eight tiles). Each tile or group of tiles can be assigned individual compression ratios. The Tier2 compression process can either be based on bit rate control or image quality.

The algorithm for this FPGA generates the rate distortion curve for each code block in each tile, performs the optimized lossy compression, and generates the JPEG 2000 byte stream.

Code Block FIFOs

The code block FIFOs are used for temporary storage of code block bytes and all or some of the JPEG 2000

byte stream. The code block bytes are moved from the Tier1 FPGA and are stored in the first two FIFOs while the Tier2 processor determines how many code block bytes will be transferred in the output byte stream. A third FIFO is either used to hold small portions of the output byte stream when using the sequential tile-by-tile format or to hold the entire output byte stream when using the tile part format. This data pathway is shown in Figure 6.

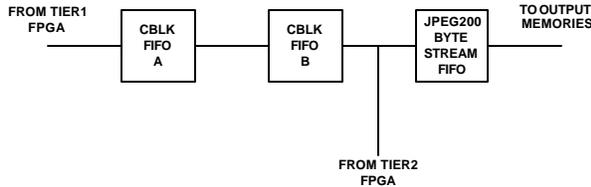


Figure 6. Code Block FIFO Block Diagram

Output Memories

Dual-port RAMs that can be read from the host computer over the PCI bus serve as the output memories for the JPEG 2000 demonstration board. The input data for these RAMs can be modified, if necessary, by an output bus from the Tier2 FPGA as shown in Figure 7. One example of such a modification is the writing of the contiguous box length into the RAM. In JP2 format, the box length is set to zero while the JPEG 2000 byte stream is stored in the last code block FIFO. This length remains zero throughout the transfer of the JPEG 2000 byte stream to the output memories. The correct length of the box is then calculated and written to the box length memory location in the output memories from the Tier2 FPGA.

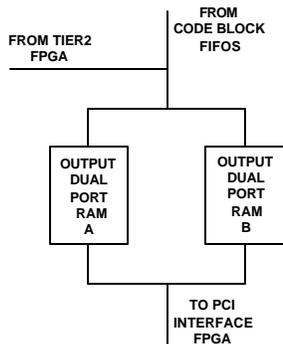


Figure 7. Output Memories

After the entire JPEG 2000 byte stream has been loaded into the output memories, it can be read from the host computer over the PCI bus. If sequential tile-by-tile formatting is used, the final byte stream is read out two

bytes at a time. If tile part formatting is used, the final byte stream is read out one byte at a time.

Image Throughput

The image throughput is determined by dividing the number of image pixels by the processing time. The processing time is defined by the period between the issuing of the start signal and the loading of the end-of-codestream marker into the last code block FIFO. The start signal is given after the image is loaded into the input memory. This timing scheme simulates a system where images are processed continuously.

APPLICATION DESIGNS

The JPEG 2000 demonstration board has provided a successful test bench for both firmware development and consideration of application designs. Figure 8 illustrates a board layout that would give the same functionality as that of the present demonstration board in a remote sensing application.

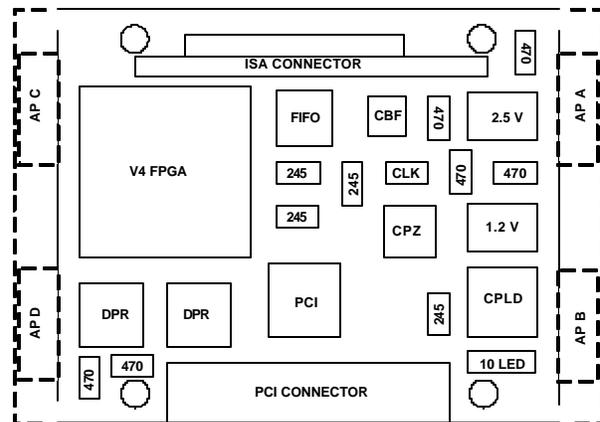


Figure 8. Small Board Design

This smaller board was designed around a PC/104 communication bus. The Xilinx Virtex 4 FPGA can hold the firmware used in both of the Virtex II FPGAs of the demonstration board, thus reducing the number of FPGAs to one. In this design, the eight input tile FIFOs were eliminated and only the input dual-port RAM was retained for converting raster images to tile format. The three code block FIFOs were reduced to one. The output dual-port RAMs remain to retain the option to convert a sequential tile-by-tile byte stream to a tile part byte stream. The other components shown in the illustration are those used for power conditioning, clocking, FPGA configuration, and host computer communications.

Depending on the needs of the user, the size of this

application board could be further reduced. For example, one of the output dual-port RAMs could be eliminated if there is no need for tile part byte stream formatting, and, if the throughput requirement is low enough, both dual-port RAMs could be eliminated.

The board shown in Figure 8 is in test form design; in this case there are four logic analyzer ports on the edge of the board, two on each side. These ports would be eliminated from a production board design, thus further reducing the size of the board. In this form, the production board would be about 4 x 4.5 inches.

CONCLUSION

The development of this board has resulted in a successful test bench for the development of JPEG 2000-based hardware systems for field applications. The knowledge gained from this project is significant in the further development of systems with faster throughput, smaller board sizes, lower power designs, and the capability to compress larger images.

References

1. Rabbani, M. and R. Joshi, "An overview of the JPEG 2000 still image compression standard," *Signal Processing: Image Communication*, vol. 17, 2002.
2. Taubman, D.S. and M.W. Marcellin, *JPEG 2000 Image Compression Fundamentals, Standards and Practice*, Kluwer Academic Publishers, 2002.
3. ISO/IEC 15444-1:2004(E), *Information technology – JPEG 2000 image coding system: Core coding system*.
4. Woolston, T.L., Holt, N., Bingham, G., and G. Wada, "A JPEG 2000 Demonstration Board," *Proceedings of the International Society for Optical Engineering*, August 2005.