FPGA-based Acceleration of Hyperspectral K-Means Clustering

David A. Kusinsky (kusinsky@ll.mit.edu)* and Miriam Leeser (mel@ece.neu.edu)** *MIT Lincoln Laboratory, Lexington, MA **Department of Electrical and Computer Engineering, Northeastern University, Boston, MA

Introduction

The goal of the current research is to investigate the use of Field Programmable Gate Arrays (FPGAs) to accelerate hyperspectral image processing algorithms. Hyperspectral image processing poses several challenges that warrant research into accelerated algorithm implementations that run on FPGA-based hardware. One such challenge can be imposed by the location of a given hyperspectral imaging (HSI) system. Airborne and space-based HSI systems rarely take advantage of high performance computing in the sensor system itself. This is due to size, weight, and power (SWAP) constraints imposed by the system platform. FPGAs are available as radiation-tolerant devices, and thus are highly reliable and preferred in space-based applications. FPGAs are also commonly used to control sensor operation and data conversion in many types of sensor systems. This places the FPGA in a unique position to perform processing of image data on a per-pixel basis. FPGAs can also provide a solid middle ground from the SWAP perspective when compared to other processing and acceleration hardware such as CPU-based computing clusters and GPU-based systems [1]. Another challenge presented in hyperspectral image processing is the large size of the image data. HSI data can be represented as a three-dimensional cube with two spatial dimensions and a third spectral dimension as shown in figure 1. This figure shows an example data cube from the Airborne Visible InfraRed Imaging Spectrometer (AVIRIS).



Figure 1: Example AVIRIS Data Cube Courtesy NASA/JPL-Caltech

This cube represents one 2-D spatial scene for each of over 200 spectral bands. This presents not only an image processing challenge, but also a storage and data communications challenge for remote platforms.

Compression of HSI data has been successfully performed using FPGAs [2]. Due to the recent advances and flight heritage of reconfigurable radiation-tolerant FPGAs, the processing of hyperspectral data using FPGAs is a low-risk option that should be considered in the development of future HSI sensor systems and image processing chains. Our initial research utilizes a host computer to provide hyperspectral data to the FPGA hardware for processing. Depending on the algorithm under consideration, processing is either completed entirely by the FPGA, or the FPGA is utilized along with the host to provide accelerated processing.

Hardware

The FPGA hardware utilized at the core of our research is a Xilinx ML605 Development Board, shown in figure 2. This board provides a Xilinx Virtex-6 FPGA as well as DDR3 memory and a PCI Express interface for host communication. Hyperspectral image data is sent to the ML605 via the PCI Express interface for processing, and the results are sent to the host when processing is completed. Host communication via the PCI Express bus is kept to a minimum whenever possible.



Figure 2: Xilinx ML605 Development Board Courtesy Xilinx

Experiments

In order to demonstrate the improved HSI data processing capabilities of the Virtex-6 FPGA, we are implementing the K-means clustering algorithm on the ML605. This algorithm was formerly implemented on a Virtex-2 FPGA by Wang [3]. For the Virtex-2 implementation, a computational speedup of 2162x was achieved, with an overall speedup of 174x when compared against a CPU implementation. The overall speedup was greatly reduced by the initial host-to-FPGA communication of the image data. One goal of our research is to improve both the computational speedup and mitigate the speedup loss caused by the host communication.

Future Work

For future work, we plan to utilize the ML605 to support other hyperspectral image processing algorithms. Some

^{*} This work is sponsored by the Department of the Air Force under Air Force Contract FA8721-05-C-0002. The opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government. Approved for public release; distribution is unlimited.

algorithm classes under consideration are dimensionality estimation, hyperspectral matched filtering [4], and spectral unmixing. These methods exist independently, and FPGA implementations for some, such as FastICA for spectral unmixing have been proposed [5]. Creation of an accelerated end-to-end processing chain (figure 3) which performs unsupervised dimensionality estimation followed by spectral unmixing or matched filtering would be beneficial to the image analysis community and is of interest for our research.



Figure 3: ML605 Processing Chain

Acknowledgements

This research makes use of software and a ML605 board donated to Northeastern University by Xilinx.

References

- A. Plaza, Q. Du and Y.-L. Chang, "High Performance Computing for Hyperspectral Image Analysis: Perspective and State-of-the-Art," *IEEE International Geoscience and Remote Sensing Symposium (IGARSS'09)*, Cape Town, South Africa, 2009
- [2] A. Plaza, "Towards Real-Time Compression of Hyperspectral Images Using Virtex-II FPGAs," *Lecture Notes in Computer Science*, Vol. 4641, pp. 248-257, Euro-Par 2007.
- [3] X. Wang and M. Leeser, "VFloat: A Variable Precision Fixed- and Floating-Point Library for Reconfigurable Hardware," ACM Transactions on Reconfigurable Technology and Systems, Vol. 3, No. 3, September 2010.
- [4] D. Manolakis, D. Marden, and G. Shaw, "Hyperspectral Image Processing for Automatic Target Detection Applications," *Lincoln Laboratory Journal*, Vol. 14, No. 1, 2003.
- [5] K. Shyu, M. Lee, Y. Wu, and P. Lee, "Implementation of Pipelined FastICA on FPGA for Real-Time Blind Source Separation," *IEEE Transactions on Neural Networks*, Vol. 19, No. 6, June 2008.