Preface

Antonio Plaza and Chein-I Chang

International Journal of High Performance Computing Applications 2008; 22; 363
DOI: 10.1177/1094342007088375

PREFACE

Antonio Plaza[1]
Chein-I Chang[2]

Hyperspectral imaging has become one of the fastest growing and emerging techniques in remote sensing science. The wealth of spectral information available from latest generation remotely sensed hyperspectral instruments (with hundreds of spectral bands in nearly continual wavelength channels) has greatly expanded the traditional *literal* (i.e. spatial domain-based) capabilities of Earth observation instruments for applications such as land use/cover classification, urban planning and management studies, environmental monitoring and assessment, and further allowed *non-literal* (i.e. spectral domain-based) capabilities in areas such as automatic target recognition for military and defense/security deployment, wildland fire detection and monitoring, biological threat detection, monitoring of oil spills and other types of chemical contamination.

Most of the above-mentioned applications require timely responses for swift decisions which depend upon high computing performance of algorithm analysis. This is particularly true for reconnaissance and surveillance applications in which moving targets are of major interest, for example vehicles in a battlefield, drug trafficking in law enforcement, or chemical/biological agent detection in bio-terrorism. Detecting and tracking these subtle targets often requires real or near real-time hyperspectral image processing algorithms. Unfortunately, because of the extremely large dimensionality of the data sets provided by hyperspectral imaging instruments, efficient hyperspectral data processing has traditionally been a major obstacle to efficient algorithm design and implementation, in particular because these algorithms often require intensive and high complexity computations to process huge volumes of data with sub-pixel precision. In future years, hyperspectral sensors will substantially increase their spatial and spectral resolution (imagers with thousands of spectral bands are currently under development) and hence hyperspectral sensors are soon expected to be able to produce a nearly continuous stream of high-dimensional data with a wealth of spatial and spectral resolution. Such an explosion in the amount, size and quality of the collected information in hyperspectral imaging presents new challenges for advanced hyperspectral data processing.

The main goal of this special issue is to provide in-depth presentations of recent research efforts addressing the need for high performance computing (HPC) for hyperspectral imaging applications. In particular, the HPC-based techniques covered by this special issue include several HPC-based architectures and techniques, including multiprocessor (e.g. commodity cluster-based) systems, large-scale and heterogeneous networks of computers, and specialized hardware architectures such as field-programmable gate arrays (FPGAs) and graphical processing units (GPUs). Combined, these topics deliver an excellent snapshot of the state of the art in those areas, and offer a thoughtful perspective of the potential and emerging challenges of applying HPC paradigms to hyperspectral imaging problems.

The special issue is composed of five articles contributed by authors with solid reputations and backgrounds in the field of hyperspectral imaging. In particular, all the techniques and methods presented in this section are well consolidated, and cover almost entirely the current spectrum of techniques and HPC architectures regularly used in hyperspectral imaging applications. We specifically avoided repetition of topics in order to provide a timely compilation of realistic and successful efforts in this field. A summary of each contribution is provided below.

The first article, "Clusters versus FPGA for Parallel Processing of Hyperspectral Imagery" by Plaza and Chang, provides a comprehensive review of the state of the art in the design of HPC systems for hyperspectral imaging. The article also presents an application case study in which a full hyperspectral data processing chain based on the pixel purity index (PPI) algorithm – included in Kodak's Research Systems ENVI, a very popular commercial software package in the remote sensing community – followed by a linear spectral unmixing technique is implemented in parallel using two different types of HPC platforms: Thunderhead, a commodity Beowulf cluster at NASA's Goddard Space Flight Center, and a Xilinx Virtex-II FPGA. Performance data are given in both cases and the advantages and disadvantages of each implementation are thoroughly discussed in the context of a real application with strong computational requirements, based

The International Journal of High Performance Computing Applications,
Volume 22, No. 4, Winter 2008, pp. 363–365
DOI: 10.1177/1094342007088375

[1]DEPARTMENT OF TECHNOLOGY OF COMPUTERS AND COMMUNICATIONS, TECHNICAL SCHOOL OF CÁCERES, UNIVERSITY OF EXTREMADURA, E-10071 CÁCERES, SPAIN (APLAZA@UNEX.ES)

[2]REMOTE SENSING SIGNAL AND IMAGE PROCESSING LABORATORY (RSSIPL), DEPARTMENT OF COMPUTER SCIENCE AND ELECTRICAL ENGINEERING, UNIVERSITY OF MARYLAND BALTIMORE COUNTY (UMBC), 1000 HILLTOP CIRCLE, BALTIMORE, MD 21250.

on the analysis of hyperspectral data collected by NASA's Jet Propulsion Laboratory over the World Trade Center area in New York City shortly after the terrorist attacks of September 11th, 2001.

The second article, "Parallel Processing of Remotely Sensed Hyperspectral Images on Heterogeneous Networks of Workstations Using HeteroMPI" by Valencia et al., discusses the role of heterogeneous parallel computing in the context of hyperspectral imaging applications. Although most available parallel processing strategies for hyperspectral image processing have traditionally assumed homogeneity in the underlying computing platform, heterogeneous networks of workstations have become a promising cost-effective solution, expected to play a major role in many on-going and planned remote sensing missions. This article develops a new parallel morphological algorithm for hyperspectral image processing using HeteroMPI, an extension of MPI for programming high-performance computations on heterogeneous platforms. In order to further analyze the performance of hyperspectral imaging applications, which typically make intensive use of many-to-one (gather) communication operations, a collective communication model is addressed. The proposed parallel algorithm is validated in the context of a hyperspectral image classification application oriented towards the classification of land-cover agricultural crops in the state of Indiana. The parallel implementations are tested using several heterogeneous clusters located at University College Dublin and a Beowulf cluster located at NASA's Goddard Space Flight Center.

The third article, "Field Programmable Gate Arrays (FPGA) for Pixel Purity Index Using Blocks of Skewers for Endmember Extraction in Hyperspectral Imagery" by Hsueh and Chang, takes advantage of a recently developed concept called blocks of skewers (BOS) to implement the PPI algorithm in real-time via FPGA design. The BOS was previously proposed to reduce computational complexity of the PPI algorithm so that it can be implemented in hardware. This article further investigates the feasibility of an efficient hardware implementation for the BOS technique in conjunction with a modified version of the PPI, referred to as BOS-APPI. Several block designs are developed to implement the BOS-APPI in hardware and thus improve the original BOS algorithm both from the viewpoint of algorithm performance and FPGA design. The implementations of different BOS designs developed in this article are implemented on a XESS Spartan II FPGA, and tested using both simulated and real hyperspectral data sets collected by the U.S. Army in the context of a military-oriented target detection application.

The fourth article, "GPU for Parallel On-Board Hyperspectral Image Processing" by Setoain et al., addresses the emerging use of GPUs for data processing aboard airborne and satellite hyperspectral imaging platforms. Driven by the ever-growing demands of the video-game industry, GPUs have evolved from expensive application-specific units into highly parallel programmable systems. In this article, several GPU-based implementations of a highly innovative morphological algorithm for hyperspectral image processing (one of the few available techniques which integrate the spatial and the spectral information in simultaneous fashion) are presented and thoroughly discussed. The proposed algorithm is implemented, using different spectral similarity metrics, on two different generations of GPUs from NVidia. The proposed GPU-based implementations are then evaluated in the context of a real mineral mapping application, using hyperspectral data sets collected by NASA Jet Propulsion Laboratory over the Cuprite mining district in Nevada.

The last article of this special issue, "Low-Complexity Principal Component Analysis for Hyperspectral Image Compression" by Du and Fowler, develops a low-complexity principal component analysis (PCA) technique to achieve computationally efficient hyperspectral data compression, an important requirement in most hyperspectral imaging missions. The PCA has been widely used in signal and image processing to reduce the dimensionality of the data to be processed. It has also become a major pre-processing technique for data compression. Unfortunately, the computational cost of implementing the PCA is generally very high; a fact that often prevents its real-time implementation. In this article, an iterative PCA (IPCA) is developed to circumvent this problem and to achieve low complexity as well as feasibility of very-large-scale integration (VLSI) implementation for real-time processing. Experimental results investigate the impacts of such low-complexity PCA on JPEG2000 compression of several hyperspectral images collected by NASA over different locations, including the Cuprite mining district in Nevada, the Jasper Ridge biological preserve in California, and the Moffett Field area at the southern end of the San Francisco Bay, California. In all cases, experiments focused both on the analysis of rate-distortion performance as well as data-analysis performance, using an anomaly-detection task for illustrative purposes.

Summarizing, the wide range of HPC platforms (including homogeneous and heterogeneous clusters and groups of clusters, large-scale distributed platforms, specialized architectures based on reconfigurable computing and commodity graphics hardware) and data processing techniques covered by the articles in this special issue exemplifies a subject area that has drawn together an eclectic collection of participants, but increasingly this is the nature of many endeavors at the cutting edge of science and technology. In this regard, one of the main purposes of this special issue is to reflect the increasing sophistication of a field that is rapidly maturing at the intersection of many different disciplines, including not only computer engineering and remote sensing, but also signal and image processing, optics, electronics, and aerospace engineering, among others. The guest editors would like to emphasize that this is the first special issue in the literature entirely devoted to providing a perspective on the state of the art of HPC techniques in the context of hyperspectral imaging problems. In order to address the need for a consolidated reference in this area, the guest editors have made significant efforts to include a collection of high-quality articles from highly recognized experts in the area. Their contributions to this special issue are greatly appreciated.

Last but not least, the guest editors would like to particularly thank the Editor-in-Chief, Prof. Jack Dongarra, for his constant support and encouragement, and to Jan Jones for all the help and support that she provided the guest editors during the process of editing this special issue. The guest editors would also like to take this opportunity to gratefully thank the reviewers who participated in the evaluation of manuscripts for the special issue for carefully assessing the content and presentation of all the contributions submitted for consideration in this special issue. Without their outstanding contributions, this endeavor could not possibly have been completed.