

The Promise of Reconfigurable Computing for Hyperspectral Imaging On-Board Systems: Review and Trends

Sebastian Lopez, *Member, IEEE*, Tanya Vladimirova, *Member, IEEE*, Carlos González, Javier
Resano, Daniel Mozos and Antonio Plaza, *Senior Member, IEEE*

Abstract— Hyperspectral imaging is an important technique in remote sensing which is characterized by high spectral resolutions. With the advent of new hyperspectral remote sensing missions and their increased temporal resolutions, the availability and dimensionality of hyperspectral data is continuously increasing. This demands fast processing solutions that can be used to compress and/or interpret hyperspectral data on board spacecraft imaging platforms in order to reduce downlink connection requirements and perform a more efficient exploitation of hyperspectral data sets in various applications. Over the last few years, reconfigurable hardware solutions such as field programmable gate arrays (FPGAs) have been consolidated as the standard choice for on-board remote sensing processing due to their smaller size, weight and power consumption when compared with other high performance computing systems, as well as to the availability of more FPGAs with increased tolerance to ionizing radiation in space. Although there have been many literature sources on the use of FPGAs in remote sensing in general and in hyperspectral remote sensing in particular, there is no specific reference discussing the state-of-the-art and future trends of applying this flexible and dynamic technology to such missions. In this work, a necessary first step in this direction is taken by providing an extensive review and discussion of the (current and future) capabilities of reconfigurable hardware and FPGAs in the context of hyperspectral remote sensing missions.

Manuscript received October 8, 2012

S. Lopez is with the Institute for Applied Microelectronics (IUMA), University of Las Palmas de Gran Canaria, Spain, e-mail: seblopez@iuma.ulpgc.es

T. Vladimirova is with the Department of Engineering, University of Leicester, UK, e-mail: t.vladimirova@le.ac.uk

C. Gonzalez is with Department of Computer Architecture and Automatics, Complutense University of Madrid, Spain, e-mail: carlosgonzalez@fdi.ucm.es

J. Resano is with Department of Computer and Systems Engineering (DIIS) at the University of Zaragoza and Engineering Research Institute of Aragon (I3A), Spain, e-mail: jresano@unizar.es

D. Mozos is with Department of Computer Architecture and Automatics, Complutense University of Madrid, Spain, e-mail: mozos@fis.ucm.es

A. Plaza is with the Hyperspectral Computing Laboratory (HyperComp), Department of Technology of Computers and Communications, University of Extremadura, Spain, e-mail: aplaza@unex.es

The review covers both technological aspects of FPGA hardware and implementation issues, providing two specific case studies in which FPGAs are successfully used to improve the compression and interpretation (through spectral unmixing concepts) of remotely sensed hyperspectral data. Based on the two considered case studies, we also highlight the major challenges to be addressed in the near future in this emerging and fast growing research area.

Index Terms— Hyperspectral remote sensing, reconfigurable hardware, field programmable gate arrays (FPGAs), hyperspectral data compression, spectral unmixing.

I. INTRODUCTION

Hyperspectral sensors are capable of generating very high-dimensional imagery through the use of sensor optics with a large number of (nearly contiguous) spectral bands, providing very detailed information about the sensed scene. From a remote sensing perspective, the spatial and significantly improved spectral resolutions provided by these latest-generation instruments has opened cutting-edge possibilities in many applications, including environmental modeling and assessment, target detection and identification for military and defense/security purposes, agriculture, urban planning and management studies, risk/hazard prevention and response including wild land fire tracking, biological threat detection, and monitoring of oil spills and other types of chemical contamination, among many others.

Because of their potential, remote sensing hyperspectral sensors have been incorporated in different satellite missions over recent years like the currently operating Hyperion on NASA's Earth Observing-1 (EO-1) satellite¹ or CHRIS sensor on the European Space Agency (ESA)'s Proba-1. Furthermore, the remote sensing hyperspectral sensors that will be allocated in future missions will have enhanced spatial, spectral and temporal resolutions, which will allow capturing more hyperspectral cubes per second with much more information per cube. For example, it has been estimated by the NASA's Jet Propulsion Laboratory (JPL) that a volume of 1 to 5 TBytes of data will be daily produced by short-term future hyperspectral missions like the NASA's HypsIRI². Similar data volume ratios are expected in European missions such as Germany's EnMAP³ or Italy's PRISMA⁴. Unfortunately, this extraordinary amount of information jeopardizes the use of these last-generation hyperspectral instruments in real-time or near real-time applications, due to the prohibitive delays in the delivery of Earth Observation payload data to ground processing facilities. In this respect, ESA have already flagged up in 2011 that "*data rates and data volumes produced by payloads continue to increase, while the available downlink bandwidth to ground stations is comparatively stable*" [1]. In this context, the design of solutions that enable to take advantage of the ever increasing dimensionality of remotely sensed hyperspectral images for real-time applications has gained a significant relevance during the last decade.

¹<http://eo1.gsfc.nasa.gov>

²<http://hyspirc.jpl.nasa.gov>

³<http://www.enmap.org>

⁴http://www.asi.it/en/flash_en/observing/prisma

Within this scenario, on-board processing systems have emerged as an attractive solution in order to decrease the delays between the acquisition of a hyperspectral image, its processing/interpretation, and the decision on a proper action to be taken according to the information extracted from the image. This can be mainly achieved in two ways: 1) performing on-board (lossless or lossy)compression of the acquired data before transmitting them, so that the remotely sensed hyperspectral images are downloaded and further processed at the ground level; and/or 2) processing the hyperspectral images according to the needs of an application (or a set of them), so that only the obtained results (i.e. number and location of thematic classes after performing a classification/clustering of the sensed images, location of a set of sought targets within an image, pure spectral signatures (*endmembers*) together with their correspondent abundance factors obtained after spectral unmixing, etc.) are transmitted. At this point, it is worth mentioning that both scenarios, which have been exemplified at Figure 1 using a toy example, are not mutually exclusive since a possible framework could be to process the image on-board and then, compress the results obtained prior to the transmission to ground.

For the particular case of Earth Observation satellites, these on-board systems should at least accomplish the following three mandatory characteristics. First, they must allow high computational performance, since all the state-of-the-art algorithms for compressing and/or processing a given hyperspectral image have a huge associated computational burden. Second, they should have compact size and reduced weight and power consumption, due to the inherent nature of remote sensing satellites. Last but not least, they must be resistant to damages or malfunctions caused by ionizing radiation, present in the harsh environment of outer space. Furthermore, it would be highly desirable that these high-performance on-board processing systems could also show a high degree of flexibility so that they can adapt to varying mission needs, faults, and/or to evolving and future processing algorithms and standards.

Among the different general-purpose high-performance computing platforms that are nowadays commercially available, current radiation-hardened and radiation-tolerant field programmable gate arrays (FPGAs) undoubtedly represent the best choice in terms of the requirements outlined above, due to their negligible size and mass when compared with traditional cluster-based systems, as well as to their lower power dissipation figures when compared with Graphics Processing Units (GPUs). Because of this reason, this work is focused on demonstrating the suitability of FPGAs to on-board processing of hyperspectral images acquired by current and future remote sensing missions, as well as on highlighting the major challenges to be addressed in the near future. More specifically, in this work we provide an extensive review of the (current and future) capabilities of FPGAs in the context of remote sensing in general and hyperspectral imaging in particular. The review covers both technological aspects of FPGA hardware and implementation issues, describing two specific case studies in which FPGAs are successfully being used to improve the compression and data interpretation (through spectral unmixing techniques) of remotely sensed hyperspectral data. We believe that this contribution is much needed; although there have been many developments in the literature discussing the

use of FPGAs in remote sensing applications, there is no specific reference discussing the state-of-the-art and future trends of applying this flexible and dynamic technology to remote sensing missions.

The remainder of this paper is organized as follows. Section II introduces the basics of FPGA technology, highlighting its most relevant characteristics for on-board systems. Section III reviews the state-of-the-art of FPGA implementations for hyperspectral imaging in general, and for compression and processing of hyperspectral images in particular. The next two sections detail these two example case studies: lossless/lossy compression (Section IV) and spectral unmixing (Section V), going from the architectural level to the FPGA implementation itself and providing implementation examples together with experimental validation and assessment using real hyperspectral scenes collected by a variety of sensors. Finally, Section VI provides some summarizing statements as well as future research hints and challenges.

II. FIELD PROGRAMMABLE GATE ARRAYS (FPGAS)

An FPGA can be roughly defined as an array of interconnected logic blocks, as it is depicted in Figure 2. One of the main advantages of these devices is that both the logic blocks and their interconnections can be (re)configured by their users as many times as needed in order to implement different combinational or sequential logic functions. In addition, modern FPGAs frequently include embedded hardware modules, such as static RAMs memories or multipliers specifically developed for digital signal processing computations, which can improve the efficiency of the system. As it is seen in Figure 3, this characteristic provides FPGAs with the advantages of both software and hardware systems in the sense that FPGAs exhibit more flexibility and shorter development times than application specific integrated circuits (ASICs) but, at the same time, are able to provide much more competent levels of performance, in terms of number of operations per watt, than general purpose processors (GPPs). In any case, these different platforms are not mutually exclusive. On the contrary, manufacturing companies such as Xilinx [2] or Altera [3] have developed System-on-a-Chip platforms that include ASICs modules, GPPs and FPGAs. These platforms can be used to develop a tightly integrated hardware/software system where the software running in the processors can take advantage of the hardware modules implemented both on the ASICs and on the FPGA. In these platforms, ASICs provide support for frequently used functions, such as those included on the communication standards, and FPGAs are used to customize the device adding hardware support specifically designed for the target applications.

FPGAs offer good performance because they can implement optimized data paths for each computational task. These datapaths take advantage of the task internal parallelism and include a customized memory hierarchy. Moreover, due to the large number of available resources in current FPGAs, frequently it is also possible to execute several tasks in parallel in order to achieve further speedups. Using optimized datapaths not only improves the performance, but also reduces the power and energy consumption when compared with GPPs. The reason is that executing a given task in GPPs involves adjusting the required computations to the

GPP instruction set and carrying out an expensive instruction decoding process, which leads to important power and energy overheads. Furthermore, the power and energy efficiency of FPGAs has significantly improved during the last decade. FPGA vendors have achieved this goal improving the FPGA architectures, including optimized hardware modules, and taking advantage of the most recent silicon technology. For instance Xilinx reports a 50% reduction in the power consumption when moving from their previous Xilinx 6 FPGAs (implemented using 40nm technology) to their most recent Xilinx 7 FPGAs (a new architecture implemented using 28 nm technology).

FPGAs are becoming an increasingly attractive solution for space-based systems not only because of their reduced size, weight, and power dissipation, as well as to their excellent tradeoff behavior between pure software and specific hardware systems, but also because the following four main reasons. First, because as with terrestrial applications, they perform well in high-throughput signal processing tasks encountered in space, like processing and/or compressing a hyperspectral image. Second, because FPGAs permit changes to the usage model and the data processing paradigm in space rather than hard-coding of all components prior to launch. In this sense, their inherent ability to change their functionality- through partial or full reconfiguration - aids in their agility, and extends the useful life of remote sensing autonomous systems. Third, because when compared to other technologies able to provide similar computational performance and design agility, FPGAs offer clear advantages in terms of size, area, and energy efficiency. At this point it is important to highlight that FPGAs have demonstrated their superiority in terms of energy efficiency with respect to other popular reduced-area low-weight kinds of general-purpose high-performance devices, such as GPUs. Fourth and last, because they can be manufactured in order to resist high levels of radiation without changing the content of their inner memories, which determines the FPGA programming, which undoubtedly constitutes another advantage of FPGAs with respect to current GPUs for space-based systems. Moreover, since FPGAs implement custom designs, additional fault-tolerant levels can be included in the system when needed, such as dual or triple modular redundancy.

Before analyzing the effects of outer space radiation in commercially available FPGAs, it is important to distinguish between antifuse, SRAM (static RAM) and Flash-based FPGAs. In antifuse-based FPGAs, the logical function to be reproduced within the targeted FPGA is obtained by burning off its internal connections, which are made of antifuses. Although these FPGAs are highly tolerant to radiation effects, their inherent nature converts them into one-time programmable devices and hence, they become less attractive for future space-based systems than SRAM and Flash-based FPGAs, since these last two can be reprogrammed. In a SRAM-based FPGA, the configuration of its logic elements and interconnections is stored in small SRAM memories distributed along the field programmable array. Since SRAM memories are volatile by default, these FPGAs must be programmed each time they are powered up, the program typically being stored in an off-chip non-volatile memory. On the contrary, in a Flash-based FPGA the configuration is stored in a non-volatile memory and hence, the FPGA device remains programmed even when the power source is removed from the system.

Besides the fact that that each way of (re)programming a FPGA has its own strengths and weakness (in this sense, interested readers are referred to [4] for a comprehensive analysis), for the scope of this work it is important to highlight that they also behave in different ways under radiation effects. These effects can be categorized into two groups: total ionizing dose (TID) effects and single event effects (SEEs). The former gives an account of the energy accumulated in the device by ionizing radiation per unit of mass, being measured in the international system of units in grays (Gy – joules per kilogram), whereas the latter define the effects caused by aisle ionizing particles, which can be destructive (in the sense that they provoke permanent damages in the device) or not. With independence of its programmability, the families of FPGAs that are up to a certain limit resistant to these radiation effects are known as radiation-tolerant and radiation-hardened FPGAs, and different companies nowadays offer such devices. For instance, Xilinx presently offers the SRAM-based Virtex-4QV[5] and Virtex-5QV [6] families, with a frequency performance up to 450 MHz, and which have been extensively analyzed in terms of radiation effects by the NASA’s Jet Propulsion Laboratory (JPL) [7] while Microsemi manufactures the RT ProASIC3 series [8], based on nonvolatile Flash-based technology, which can run up to 350 MHz.

Different state-of-the-art works have demonstrated that both families of reprogrammable FPGAs are suitable for space-based systems, although SRAM-based FPGAs are in general more tolerant than Flash-based FPGAs to TID effects, while on the contrary, SRAM-based FPGAs are more vulnerable than Flash-based FPGAs to SEEs [9]. In the remainder of this paper we will center our efforts in demonstrating that both types of FPGAs can be considered for future on-board hyperspectral imaging systems by means of two different case studies: compression and linear unmixing.

III. REVIEW OF STATE-OF-THE-ART FPGA-BASED HYPERSPECTRAL IMAGING SYSTEMS

Recently, significant efforts have been made directed towards the increase of the performance of remote sensing applications, specifically in the acceleration of hyperspectral imaging algorithms [10-14]. In the field of hyperspectral remote sensing, mainly four solutions have been addressed: cluster computing, heterogeneous computing, GPUs, and FPGAs. The motivation for this effort relies on the high computation requirements needed for achieving real or near real-time processing capabilities. Due to the superiority of FPGA devices for on-board systems, this section is focused on reviewing the state-of-the-art works that deal with the implementation of hyperspectral imaging algorithms onto FPGAs. More precisely, this section has been divided into three subsections devoted to review the existing FPGA-based implementations for hyperspectral image compression, spectral unmixing, and other hyperspectral imaging algorithms.

A. State-of-the-art hyperspectral image compression implementations on FPGAs

Due to the large amount of data generated by sensors, especially spaceborne ones, it is necessary to reduce the size of the data in order to download all the acquired sensor input. Image compression compensates for the limited on-board resources, in terms of

mass memory and downlink bandwidth and thus provides a solution to the “*Bandwidth Versus Data Volume*” dilemma of modern spacecrafts. The problem becomes increasingly important for hyperspectral remote sensing systems with the current trend towards developing more accurate hyperspectral sensors covering hundreds of spectral bands. Thus, hyperspectral image compression becomes an important on-board capability in order to alleviate the memory costs and communication bottleneck in present and future satellite missions [15].

1) Introduction to compression of satellite imagery

Image compression methods can be divided into two classes, lossless or lossy. With lossless image compression, the reconstructed image is exactly the same as the original one, without any information loss at the expense of achieving relatively modest compression ratios when compared with lossy methods. This is because the entropy, which measures the quantity of information contained in a source, imposes a theoretical boundary on the lossless compression performance, expressed by the lowest compression bit-rate per pixel. Entropy depends on the statistical nature of the source and ideally an infinite-order probability model is needed to evaluate it. On the contrary, lossy image compression enables competitive compression ratios at the expense of introducing a varying degree of unrecoverable information loss in the reconstructed images.

The process of lossless compression is achieved via removing the redundancy in the data. There are several types of redundancies in a remotely sensed image, such as spatial redundancy, statistical redundancy, human vision redundancy and spectral redundancy.

Spatial or intra-band redundancy means that the pixel information could be partially deduced by neighboring pixels. Spatial decorrelation methods, like prediction or transformation, are usually employed to remove the spatial redundancy. Prediction is used to predict the current pixel value from neighboring pixels. For example, the Differential Pulse Code Modulation (DPCM) method is a typical prediction based technique. On the other hand, transformation is used to transform the image from the spatial domain into another domain, applying, for example, the Discrete Cosine Transform (DCT) or the Discrete Wavelet Transform (DWT) [15].

Statistical redundancy explores the probability of symbols. The basic idea is to assign short codewords to high probability symbols, and long codewords to low probability symbols. Huffman or Arithmetic coding are two popular methods to remove statistical redundancy which are usually known as entropy coding methods. Human vision redundancy explores the fact that eyes are not so sensitive to high frequencies. Removing human vision redundancy is normally achieved by quantization, with high frequency elements being over quantized or even deleted. Spectral or inter-band redundancy is present in three dimensional (3-D) hyperspectral imagery with a large number of spectral bands in addition to the spatial redundancy present in two dimensional (2-D) digital images, adding a third dimension to compressing 2-D images.

A typical architecture of a 2-D image compression system consists of a spatial decorrelation stage, which is followed by a

quantization and entropy coding stages, exploiting the first three redundancies described above. Based on techniques used for spatial decorrelation, compression systems can be divided into prediction, DCT and DWT-based systems. Prediction-based compression methods used in space missions include DPCM [16], [17], Adaptive DPCM [18], the CCSDS Lossless Data Compression (CCSDS-LDC) [19], Lossless JPEG [20] and JPEG-LS [21]. DCT-based compression methods include JPEG-baseline [20] and specifically designed DCT compression methods. DWT-based compression methods include JPEG2000 [22], Embedded Zerotree Wavelet (EZW) [23], SPIHT [24], CCSDS Image Data Compression (CCSDS-IDC) [25] and specifically designed DWT compression methods.

The Consultative Committee for Space Data Systems has played and will continue playing a significant role in providing efficient on-board compression methods. In May 1997, CCSDS published a recommendation standard for lossless data compression, which is an extended Rice algorithm with added two low-entropy coding options [19]. This recommendation addresses only lossless source coding and is applicable to a large variety of digital data, generated by different types of imaging or non-imaging instruments on board satellites. The algorithm consists of two separate functional parts: a preprocessor and an adaptive entropy coder. The pre-processor is used to de-correlate a block of J sample data and subsequently map them into symbols suitable for the entropy coding stage. The entropy coding module is a collection of variable-length codes operating in parallel on blocks of J preprocessed samples. Each code is nearly optimal for a particular geometrically distributed source. The coding option achieving the highest compression is selected for transmission, along with an ID bit pattern used to identify the option to the decoder. Because a new compression option can be selected for each block, the algorithm can adapt to changing source statistics.

Since 1998, the CCSDS data compression working group began to assess the feasibility of establishing an image compression recommendation suitable for space applications and the CCSDS-IDC Blue Book was finally produced in November 2005 [25]. The compression technique described in this recommendation can be used to produce both lossy and lossless compression. It supports both frame-based input formats produced, for example, by charge-coupled device (CCD) arrays and strip-based input formats produced by push-broom type sensors. An image pixel resolution of up to 16 bits is supported. The compressor consists of two functional parts, a DWT module that performs decorrelation and a Bit-Plane Encoder (BPE), which encodes the decorrelated data. Although similar to that of JPEG2000 standard, this architecture differs from it in several respects: a) it specifically targets high-rate instruments used on board space missions; b) a trade-off is performed between compression performance and complexity with a particular emphasis on space applications; c) the lower computational complexity of the CCSDS-IDC algorithm supports a fast and low-power hardware implementation; d) it has a limited set of options, enabling its successful application without an in-depth algorithm knowledge. According to literature sources CCSDS-IDC could achieve performance similar to that of JPEG2000 [26-29].

2) *Basics of hyperspectral image compression*

When compressing hyperspectral images, both the spectral and spatial types of redundancy need to be removed in order to achieve a good compression performance via a lossless or a lossy process. In most of the cases, the spectral decorrelation is performed first, followed by the spatial decorrelation, which is used in 2-D image compression, as depicted in Fig. 4. The spectral decorrelation aims to reduce the spectral redundancy that exists between bands, whereas spatial decorrelation takes care of the inter-pixel redundancy within a band. This scheme has been used widely in remote sensing as well as in medical applications for three dimensional medical data.

Each of the spectral and spatial decorrelation processes can be performed by either a lossless or a lossy algorithm as shown in Fig. 4. To achieve an overall lossless compression process, both the spectral and the spatial decorrelation stages should utilize lossless transformations. On the other hand, an overall lossy compression process can utilize not only lossy spectral and lossy spatial decorrelation modules but also lossless spectral decorrelation and lossy spatial decorrelation modules and vice versa.

Different methods have been proposed for hyperspectral data compression, such as (i) predictive coding, e.g. Differential Pulse Code Modulation, CCSDS Lossless Multispectral & Hyperspectral Image Compression standard; (ii) Vector Quantization (VQ); (iii) transform coding, e.g. Karhunen-Loève Transform (KLT), DCT, DWT. VQ methods can be seen as coding of the hyperspectral image in a cube form, where the spatial and spectral decorrelations are processed in one single stage, while predictive and transform coding methods have been typically used to tackle both spectral and spatial decorrelation [30].

The early hyperspectral image compression studies were based on DPCM. This predictive coding technique predicts a current pixel value using the neighboring pixels and makes use of the difference between the real and predicted values. DPCM can be employed for spatial, spectral and spectral-spatial decorrelation. Spatial predictive methods have been upgraded to perform inter-band compression via increasing the size of the neighborhood from 2-D to 3-D. However, according to [31] the direct extension from 2-D to 3-D may not always provide tangible benefits, and sometimes can prove to be detrimental. Therefore, it is necessary to develop predictors that are specialized for 3-D hyperspectral images.

Vector quantization is a form of pattern recognition, where an input pattern (i.e. the hyperspectral image) is ‘approximated’ by a predetermined set of standard patterns. The set of standard patterns is also known as a codebook. The difference between the original and the approximated data and the codebook address are the compressed data that is needed for decompression. Examples that use VQ for hyperspectral image compression are given in [32], [33]. VQ differs from the compression scheme described in Fig. 4, since it encodes the data in one stage, covering both the spatial and the spectral domains. However, in [34] DCT was applied in the spectral domain in order to compress the residual data produced by the mean-normalized vector quantization (M-NVQ) algorithm in the spatial domain, which may be seen as compliant with Fig. 4.

In transform coding, the original hyperspectral data is projected by a set of basis vectors to produce a set of product values. The basis vectors differ depending on the particular transform used, e.g. DWT, DCT or KLT. Transform coding such as DWT and DCT can be used for either spectral or spatial decorrelation in hyperspectral compression. KLT, on the other hand, has been used for spectral decorrelation due to its intrinsic energy-compacting capability that is based on a statistical method. For example in [35] spectral decorrelation is performed with KLT and spatial decorrelation with DCT, in [36] both, spectral and spatial decorrelation are performed with DWT and in [37] spectral decorrelation is performed with Integer KLT and spatial decorrelation with DWT.

At this point, it should be highlighted that CCSDS have recently published a new recommended standard for lossless multispectral and hyperspectral image compression for space applications [38]. It is based on a predictive coding method that depends on the values of nearby samples in the current spectral band (where the reference pixel, $S_{z,y,x}$ is located) and P preceding spectral bands, where P is a user-specified parameter. In each spectral band, a local sum of neighbouring sample values is calculated through neighbour-oriented or column-oriented approach, which is used to compute one or more local differences within the spectral band. A predicted sample value, $\hat{S}_{z,y,x}$ is then calculated by using an adaptive-weighted sum of the local differences in the current and P previous spectral bands. Following that the difference between $S_{z,y,x}$ and $\hat{S}_{z,y,x}$ is mapped to an unsigned integer $\delta_{z,y,x}$ to produce the predictor output. Finally, the predictor output is coded using an entropy coder that is adaptively adjusted to adapt changes in the statistics of the mapped prediction residuals as defined in [39]. The algorithm provides spatial-spectral decorrelation in one stage.

Finally, it is worth noting that very few remote sensing missions have included hyperspectral image compression capabilities on board so far. Among the missions that are operational at present only three spacecraft - EO-1, Mars-Express and IMS-1 - perform hyperspectral image compression on board. However, details about the used compression techniques and their implementation are not available in the open literature. The importance of on-board hyperspectral image compression will grow in the future, as it is expected that newly developed hyperspectral instruments will provide greater spectral coverage, generating tremendous amounts of valuable data [30].

3) *FPGA-based hyperspectral image compression systems*

Due to their critical role in hyperspectral imaging systems, compression algorithms have been mainly implemented on FPGAs for on-board exploitation. For example, in [40] a Xilinx XC3S4000 FPGA is used to accelerate the critical path of a lossless hyperspectral image compression algorithm. A speedup of 21x is achieved when compared with the software version with a processing capability of 16.5 megapixels per second. The compressor proposed in [41] is also lossless in nature and has been implemented using a Xilinx Virtex-4 LX25 FPGA, resulting in an occupation of 43% and a power consumption of 1.27 watts

(W). Compared with the software version the speedup factor was 58x with a maximum data transfer of 58 megasamples per second. In [42], authors have presented a system based on the consultative committee for space data systems (CCSDS) recommendation, capable both of lossless and lossy image compression using three FPGAs from Xilinx: a Spartan3E 12000E-5, a Virtex-4 LX25-12 and a Virtex 5-LX30-3. Results are provided for three different optimization cases: case 1, with no optimizations; case 2, with supreme quantization; and case 3, with supreme quantization and multiplier free. For the best case, a throughput of 250 Mpixels/s at 8 bits/pixel was achieved. Depending on the optimizations performed, the power consumption on the Spartan3E was of 309 mW for case 1, 1286 mW for case 2 and 273 mW for case 3, demonstrating that optimization efforts have a positive effect in the power consumption. In [43] a new compressor for hyperspectral images is presented. The compressor uses a linear prediction between bands to exploit the inter-band correlation. Although the authors claim that an FPGA implementation has been developed, only a rough estimation based on other references is provided. Another approach for compression is followed in [44], where the author uses the concept of spectral unmixing to extract the endmembers and abundances and express the remainder pixels as a linear combination of the extracted endmembers. For compression ratios of 20:1, 40:1 and 80:1, the spectral angle similarity scores among the USGS reference signatures from a well-known hyperspectral image collected over the Cuprite mining district in Nevada by the airborne visible infra-red imaging spectrometer (AVIRIS)⁹ show that the proposed compression system achieves better quality than JPEG2000 multi-component [22] or SPIHT (set partitioning in hierarchical trees) [45]. The compression algorithm is implemented on a Virtex-II XC2V6000-6 FPGA [46], using only 36% of the resources and when compared against a sequential implementation using a PC running at 2.6 GHz, the speedup factor was of 70x, using 400 parallel processors and a transfer rate of 40 Mbytes/second. Although the results are obtained very fast (7.94 seconds for the whole compression procedure), the response is not strictly real-time as the AVIRIS instruments must process a full image cube (614×512 pixels with 224 bands) in no more than 5 seconds to fully achieve real-time performance. This same compression scheme has been also implemented in a GPU from NVidia, the C1060 Tesla GPU, and compared with a FPGA implementation, this time using the Virtex-II PRO xc2vp30 from Xilinx [47]. The total processing time for compressing an AVIRIS hyperspectral scene collected over the World Trade Centre in Manhattan was of 31.23 seconds using the FPGA and utilizing approximately 76% of the resources at a frequency of 187 MHz. Using the GPU, the total processing time was of 17.59 seconds. Compared with the sequential processing on an Intel Core i7 920, with a total processing time of 1073.03 seconds, the speedup factor of the FPGA implementation was of 34.52x and the speedup factor of the GPU was of 61.28x. Even taking into account that the GPU exhibits a speedup factor of 1.77x compared with the FPGA, the power dissipation of this FPGA is of about 221.85mW [48], while the typical power consumption of this particular GPU is about 200W. This means that the GPU consumes 901 times more power than the FPGA, achieving only 1.77 times more speed than the FPGA.

⁹ <http://aviris.jpl.nasa.gov>

B. State-of-the-art hyperspectral linear unmixing implementations on FPGAs

The number and variety of information extraction tasks in hyperspectral remote sensing is enormous [49]. However, one of the main problems in hyperspectral data exploitation is the presence of mixed pixels, which arise when the spatial resolution of the sensor is not enough to separate spectrally distinct materials. Furthermore, no matter what the spatial resolution is, the spectral signatures collected in natural environments are invariably a mixture of the signatures of the various materials found within the spatial extent of the ground instantaneous field view of the imaging instrument. Within this context, spectral unmixing has rapidly become one of the most popular techniques to analyze hyperspectral data, which allows us for sub-pixel characterization [50].

1) Basics of hyperspectral linear unmixing

Spectral unmixing generally comprises two stages: (i) identification of pure spectral signatures (endmembers) in the data, and (ii) estimation of the abundance of each endmember in each (possibly mixed) pixel [50]. A standard technique for spectral mixture analysis is linear spectral unmixing, which assumes that the collected spectra at the spectrometer can be expressed in the form of a linear combination of endmembers weighted by their corresponding abundances. It should be noted that the linear mixture model assumes that the secondary reflections and the scattering effects can be neglected in the data collection procedure, and hence the measured spectra can be expressed as a linear combination of the spectral signatures of materials present in the mixed pixel. If the impact of the secondary reflections or the scattering effects is relevant, more complex non-linear models can be applied but they demand a priori information about the geometry and physical properties of the observed objects and also increase the computational complexity [49]. In the following, we will focus on the linear model since it is the most commonly applied for hyperspectral unmixing.

To define the linear spectral unmixing problem in mathematical terms, let us assume that a remotely sensed hyperspectral scene with n bands is denoted by I , in which the pixel at the discrete spatial coordinates (i, j) of the scene is represented by a vector $X(i, j) = [x_1(i, j), x_2(i, j), \dots, x_n(i, j)] \in \mathbb{R}^n$, where \mathbb{R} denotes the set of real numbers in which the pixel's spectral response $x_k(i, j)$ at sensor channels $k = 1, \dots, n$ is included. Under the linear mixture model assumption, each pixel vector in the original scene can be modeled using the following expression:

$$x(i, j) \approx \sum_{j=1}^p e_j \cdot \Phi_j(i, j) + w = E \cdot \Phi + w \quad (1)$$

where e_j denotes the spectral response of an endmember, Φ_j is a scalar value designating the fractional abundance of the endmember e_j , p is the total number of endmembers, and w is a noise vector. It should be noted that $E = \{e_j\}_{j=1}^{j=p}$ can be seen as an $n \times p$ matrix. The solution of the linear spectral mixture problem described in (1) relies on the correct determination of a set of p

endmembers and their correspondent abundance fractions $\Phi = \{\Phi_j\}_{j=1}^{j=p}$ at each pixel x .

Over the last decade, several algorithms have been developed for automatic or semiautomatic extraction of spectral endmembers [50], [51]. Winter's N-FINDR [56] has been one of the most successfully applied techniques for automatically determining endmembers in hyperspectral image data. The algorithm attempts to automatically find the simplex of maximum volume that can be inscribed within the hyperspectral data set. The procedure begins with a random initial selection of pixels [see Fig. 6(a)]. Every pixel in the image must be evaluated to refine the estimate of endmembers, looking for the set of pixels that maximizes the volume of the simplex defined by selected endmembers. The corresponding volume is calculated for every pixel in each endmember position by replacing that endmember and finding the resulting volume [see Fig. 6(b)]. If the replacement results in an increase of volume, the pixel replaces the endmember. This procedure is repeated until there are no more endmember replacements [see Fig. 6(c)].

On the other hand, the image space reconstruction algorithm (ISRA) [57] is one of the most popular algorithms for abundance estimation in hyperspectral image data, including expectation-maximization maximum likelihood (EMML) [52], fully constrained least-squares unmixing (FCLSU) [53], and non-negative constrained least-squares unmixing (NNLSU) [54]. However, these algorithms are computationally intensive and place a heavy burden on computing systems, and hence, they demand efficient hardware for scenarios under tight time constraints.

2) *FPGA-based hyperspectral unmixing systems*

In [55] an implementation of the N-FINDR algorithm using a Virtex-4 XC4VFX60 FPGA from Xilinx was developed. This FPGA model is similar to radiation-hardened FPGAs certified for space operation. The experimental results show that the hardware version of the N-FINDR [56] algorithm can significantly outperform an equivalent software version while being able to provide accurate results in near real-time. The speedup of this implementation, compared with a software description developed in C language and executed on a PC with AMD Athlon 2.6 GHz processor and 512 Mb of RAM, is 37.29x for AVIRIS Cuprite (16 endmembers), 38.10x for a hyperspectral image collected also in the Cuprite mining district by EO-1 Hyperion (21 endmembers), and 37.70x for an AVIRIS image collected over the Jasper Ridge biological preserve in California (19 endmembers). This average speedup factor of 37.63x is quite constant across all the images, even taking into account the differences in the number of endmembers. The second main problem, together with the extraction of suitable endmembers, is the estimation of the fractional abundances of endmembers in each pixel of the hyperspectral scene. An FPGA implementation of the abundances computation task using a parallel image space reconstruction algorithm (ISRA) [57] is described in [58]. The FPGA was the same used in [55], i.e. the Virtex-4 XC4VFX60 FPGA from Xilinx. The system includes a direct memory access (DMA) module and implements a pre-fetching technique to hide the latency of the input/output communications, one of the main bottlenecks found in this kind of applications. In [58] the number of ISRA modules used in parallel is 16, achieving a speedup

factor of 10x when processing the AVIRIS Cuprite scene and over 12x when it comes to the two Jasper Ridge AVIRIS scenes. Authors also reach the conclusion that, using FPGAs, the execution time scales linearly with the size of the image. On the other hand, the software implementation increases 3.1 times, which is clearly worse than the hardware implementation. In [59] and [60], an FPGA implementation of the pixel purity index (PPI) algorithm [61] for endmember extraction using the Virtex-II PRO xc2vp30 is presented. The proposed hardware system is easily scalable and able to provide accurate results with compact size in near real-time, which makes the reconfigurable system suitable for on-board hyperspectral data processing. Results show that, against the serial software version which takes 3068 seconds to process the whole image, the proposed FPGA implementation takes only 31 seconds, using 10000 skewers and launching 100 skewers in parallel, achieving a speed-up factor of 98.96x. Another FPGA implementation of the PPI algorithm is presented in [62], in this case the FPGA needs 62 seconds to process the same image, hence the speedup factor is smaller (49.48x).

In [63], ISRA has been used in hyperspectral imaging applications to monitor changes in the environment and specifically, changes in coral reef, mangrove, and sand in coastal areas. In particular, authors have faced the problem using a hardware/software co-design methodology. The hardware units were implemented on a Xilinx Virtex-II Pro XC2VP30 FPGA and the software was implemented on the Xilinx Microblaze soft processor. As it has been observed in all the previous references, the main bottleneck found in this implementation was again data transfer. The only implementation data provided in this paper is that the FPGA was divided in three components: numerator, denominator, and multiplier, where each component works at operating frequencies of 93.75 MHz, 84.99 MHz and 113.14 MHz, respectively.

Finally, we would like to highlight two very recent works that deal with the FPGA implementation of two different algorithms for hyperspectral endmember extraction. The first one [64] proposes an architecture for implementing onto a generic FPGA device the so-called real-time fast simplex growing algorithm (RT-FSGA), which is derived from the simplex growing algorithm (SGA) [65], [66], together with the fast computation of simplex volumes uncovered in [67]. One of the main advantages of this architecture comes from the fact that it allows the number of endmembers (p) to vary with image data sets to accommodate various values as opposed to being fixed in the N-FINDR published in [55]. Unfortunately, results about the logic resources occupied by the proposed architecture and/or about its maximum running frequency within a FPGA are not available since the authors have not mapped their architecture onto a FPGA, or at least, they have not disclosed their results. In this sense, only results concerning the speedup of an hypothetical implementation onto a generic FPGA working at 50 MHz with respect to a MATLAB description running onto an Intel Core2 Quad CPU@2.5 GHz and 2-GB memory are reported, resulting in a factor that ranges from 229 to 456 depending on the input image. The second of these works [68] proposes a novel FPGA-based architecture implementing the modified vertex component analysis (MVCA) algorithm [69]. In particular, two versions of the MVCA algorithm which differ on the use of floating point or integer arithmetic for iteratively projecting the hyperspectral cube

onto a direction orthogonal to the subspace spanned by the endmembers already computed were mapped onto a XC5VSX95T FPGA from Xilinx, which is quite similar to the new generation of radiation hardened reconfigurable FPGAs from the same company (Virtex-5QV series). With respect to the percentage of the FPGA resources occupied by both versions of the proposed architecture, the authors report that the number of slice registers used in the floating point implementation varies from 27% ($p=3$) to 80% ($p=15$) and the number of slice look-up tables (LUTs) varies from 21% ($p=3$) to 61% ($p=15$), while for the case of the integer implementation, the number of slice registers used varies from the 27% ($p=3$) to the 71% ($p=15$) while the number of slice LUTs varies from 21% ($p=3$) to 51% ($p=15$). Moreover, all the synthesized integer precision architectures can operate with a frequency up to 268.152 MHz while the maximum frequency achieved for the synthesized floating point architectures has been 210.438 MHz, being these maximum working frequencies independent of the number of endmembers to be extracted thanks to design strategy followed by the authors. These results demonstrate that the FPGA implementation of the integer version of the MVCA algorithm shows a better performance in terms of hardware resources and processing speed than its floating point counterpart, both of them being capable of processing hyperspectral images captured by the NASA's AVIRIS sensor in real-time once they are loaded in the internal FPGA memories.

C. Other state-of-the-art hyperspectral imaging implementations on FPGAs

Several other hyperspectral imaging algorithms have been implemented in FPGAs for improved exploitation. In [70], an implementation of independent component analysis (ICA) [71] is made in order to reduce the dimensionality of hyperspectral images. A parallel ICA has been implemented on a Virtex V1000E running at a frequency of 20.161MHz, and the board transfers data directly with CPU on the 64-bit memory bus is made at the maximum frequency of 133MHz. Unfortunately, authors do not provide execution times and comparisons with other similar ICA implementations ([72], [73]) are made based only on maximum synthesized frequencies and the size of the observation data sets.

Another important field of application is hyperspectral image classification. The work in [74] explores design strategies and mappings of classification algorithms for a mix of processing paradigms on an advanced space computing system, featuring MPI-based parallel processing with multiple PowerPC microprocessors, each coupled with kernel acceleration via FPGA and/or AltiVec resources (in-chip vector processing engine). The case-study algorithm adopted for this research uses a linearly constrained minimum variance (LCMV) beam-forming approach, as described in [75], and has a similar computational structure to many other algorithms. Design of key components for hyperspectral imaging systems as the autocorrelation matrix calculation, weight computation, and target detection are discussed, while the hardware/software performance tradeoffs are evaluated. Several parallel-partitioning strategies are considered for extending single-node performance to a clustered architecture consisted on a 10-node PowerPC cluster. Each node contained a 1.4 GHz PowerPC 7455 with AltiVec and 1GB of SDRAM, connected together with Gigabit Ethernet. Additionally, four nodes were equipped with ADM-XRC-4 FPGA boards from Alpha Data, each

containing a Xilinx Virtex-4 SX55 FPGA and four independent 4MB SRAM modules. Speedup factors compared with the serial simple software version and using AltiVec and 64, 256 and 1024 bands were of 6.47x, 7.31x and 5.41x respectively, applying the classification algorithm to three 512x512 images. The highest speedup took place for 256 bands. When using FPGAs, the speedup factors were of 2.27x, 7.48x and 14.23x for the same amount of bands. It is clear that the use of FPGA is meaningful in the case of high-dimensional images in the spectral domain. When using 4 processing nodes, the speedup factor using AltiVec is 3.46x, with an average power consumption of 178.75 W, including the CPU and a total energy consumption of 101.5 kJ. For the same amount of nodes, the FPGA accelerated version has a speedup of 32.86x, an average power consumption of 172.51 W and a total energy consumption of 10.3 kJ. Hence, it is clear how the FPGA based systems are capable of higher execution speedups with lower power consumptions.

Another important hyperspectral application which requires on-board processing is target detection. In the literature, several efforts have been discussed towards the efficient implementation of target detection algorithms in FPGAs. For instance, in [11] several examples of FPGA implementations for target detection algorithms are discussed. Specifically, chapter 15 in [11] generally discusses the use of FPGAs in detection applications and provides specific application case studies in cloud detection and dimensionality reduction for hyperspectral imaging. Chapter 16 in [11] describes an FPGA implementation of constrained energy minimization (CEM), which has been widely used for hyperspectral target detection applications [76]. The implementation is carried out using the coordinate rotation digital computer (CORDIC) algorithm [76] to convert a Givens rotation of a vector to a set of shift-add operations. This strategy allows for efficient implementation of the CEM in FPGA architectures. Chapter 17 in [11] describes an on-board real-time processing technique for fast and accurate target detection and discrimination in hyperspectral data. In [77], a target detection system based on constraint linear discriminant analysis (CLDA) [78] is implemented using FPGAs. The system is aimed at real-time or near real-time target detection using hyperspectral data. Although the paper only provides resources consumed by the FPGA and total cycles for the matrix inversion, authors claim that the developed hardware accelerator is able to achieve real-time or near real-time throughput depending on the input pixel vector size. Recently, [79] presented an FPGA design for efficient implementation of the automatic target detection and classification algorithm (ATDCA) [76] on two different kinds of FPGA architectures: Xilinx Virtex-5 and Altera Stratix-III. The experimental results indicate significant performance gains in parallel execution of the ATDCA algorithm but not with real-time performance. In [80], a real-time target detection architecture for hyperspectral image processing is presented and discussed. The proposed architecture is designed and implemented in FPGA to illustrate the relationship between hardware complexity and execution throughput of hyperspectral image processing for target and anomaly detection [76], obtaining relevant results in terms of real-time performance in the context of different applications. In [81], real-time implementations of several popular detection and classification algorithms for hyperspectral imagery are discussed. The adopted approach to real-time implementation of such

algorithms is based on using a small portion of hyperspectral pixel vectors in the evaluation of data statistics. An empirical rule of an appropriate percentage of pixels to be used is investigated, which results in reduced computational complexity and simplified hardware implementation. Last but not least, [82] recently discussed the problem of implementing matched filters on FPGAs, the Cell IBM processor and GPUs. The matched filter is an important technique for the processing of hyperspectral data, particularly in the context of target detection applications. In this work, the performance of a matched filter algorithm implementation on an FPGA-accelerated co-processor (Cray XD-1), the IBM Cell microprocessor, and the NVIDIA GeForce 7900 GTX GPU graphics card has been evaluated. Authors implement only a reduced set of operations for the matched filter, although it is fair to mention that the most time consuming parts have been the ones considered. In all the cases, the CPU, bus interfaces, power supplies, and involved support hardware are included in all of the performance measurements, including power. The speedup factor (compared with the serial CPU version) of the FPGA was of 3.91x with a power dissipation of 350 W. The speedup factor of the GPU version was of 3.1x, with a power dissipation of 350 W. Finally, the speedup factor of the Cell processor was of 8x, with a power dissipation of 315 W. In this case, the best performance, considering speedup and power, is achieved by the Cell IBM processor, mainly due to its native vector processing arithmetic, which easily maps the matched filter modules.

IV. FIRST CASE STUDY: COMPRESSION OF HYPERSPECTRAL IMAGES

In this section, we present an image compression system based on the Karhunen-Loève transform (KLT) as a case study. In order to provide a better understanding, the contents of this section have been divided into three subsections. Subsection IV.A introduces the KLT and its reversible version, known as Integer KLT. Subsection IV.B details the structure of the spectral decorrelation modules based on the KLT and Integer KLT transforms proposed in this works and, finally, Subsection IV.C discloses the most significant results achieved.

A. *The KLT and Integer KLT transforms*

Spectral decorrelation is a unique component in hyperspectral compression systems, while spatial decorrelation is similar to 2-D image compression performed in other applications. This section describes the results of a research effort, aimed at the development of lossless and lossy spectral decorrelators for satellite hyperspectral image compression using KLT. The work enables the on-board implementation of all hyperspectral image compression options, shown in Fig.4, assuming that existing spatial decorrelation solutions are employed, such as the one reported in [42].

The Karhunen-Loève transform is an orthogonal linear transform, which is applied to 3-D data sets to decorrelate data across different bands, leading to a more compressible data set. KLT, which is strongly related to the well-known Principal Component Analysis (PCA) technique, has been used for multi-component image compression, not only in remote sensing applications but also in magnetic resonance imaging (MRI), facial recognition, as a technique for finding patterns in high dimensional data, etc..

KLT is considered the optimum method to spectrally decorrelate multispectral data with a number of spectral bands above four according to [35]. The KLT algorithm was reported as the best spectral decorrelator for hyperspectral image compression in [83] where JPEG2000 was used as the spatial compression method.

The reversible version of KLT, which is called the Integer KLT, performs lossless spectral decorrelation. Experimental results have confirmed that the Integer KLT outperforms other techniques in terms of compression ratio when applied to lossless hyperspectral compression [84].

However, the highly computationally intensive nature of KLT and Integer KLT is a major challenge when implementing these algorithms. In addition, KLT does not have a fast computation scheme, unlike other transforms, such as DCT and the discrete Fourier transform (DFT). The KLT computation flow consists of the following processes [85]:

- BandMean - obtains the mean of each band
- MeanSub - subtracts each band from its mean
- Covariance Matrix - obtains the covariance matrix of the MeanSub
- Eigenvectors - obtains the eigenvectors of the covariance matrix, which represent the principle components of the image
- Multiplication of the eigenvectors by the MeanSub

While some of the computation processes above are simple repetitive operations, such as BandMean and MeanSub, others, such as the covariance matrix and the eigenvector evaluations, are more complicated processes involving various sequential operations. Table I summarizes the KLT computational requirements, where M and L represent the spatial coordinates and n - the number of bands of the hyperspectral image, represented as a 3-D array of $M \times L \times n$. From Table I, it can be seen that the computational intensity is proportional to the dimensions of the image. It can also be concluded that the number of the multiplication, addition and subtraction operations is significantly higher than that of the division and trigonometric operations [85].

A distinctive feature of the KLT algorithm, when it is employed in hyperspectral image compression, is that the size of the processed matrix in the eigenvectors computations is very high, since its dimensions are equal to the number of the spectral bands, which are usually in the order of 100s. The resultant computational complexity can be overcome through a clustering approach to KLT, which is found to reduce the processing time and memory requirements although it compromises the compression ratio [86]. The negative effect on the compression ratio can be minimized by increasing the size of the clusters. In addition, it is found that selecting the cluster size to be equal to a power of 2 (i.e. 2, 4, 8, 16, 32 etc.) speeds up the multiplications in the KLT Covariance and BandMean modules [87].

The Integer KLT, proposed in [88], which represents the output in an integer form, is an approximation of KLT, based on matrix factorization. Similarly to KLT, the processes involved in the Integer KLT computation are performed on large matrices and

therefore they are computationally intensive, which slows down the Integer KLT evaluation significantly. In addition to BandMean, covariance matrix and eigenvectors computations as in the original KLT, the Integer KLT includes two more complex sequential processes: matrix factorisations and lifting. As in KLT, the compression of a hyperspectral image with n number of bands will involve generating an eigenvector matrix A of size $n \times n$ from the covariance matrix between each pair of bands. Matrix factorization will be applied on the A matrix, which is a non-singular matrix into four $n \times n$ matrices: a permutation matrix, P , and three other matrices called triangular elementary reversible matrices (TERMs): L (lower TERM), U (upper TERM) and S (lower TERM). The factorization is not unique and depends on the pivoting method used that will affect the error between the integer approximation and the original KLT transformation. The intrinsic energy-compacting capability of KLT will be affected by the factorization, so the error should be minimized as much as possible.

B. Hardware designs

Although hardware KLT implementations have been proposed previously, very few authors have targeted embedded computing platforms, such as [89], where a parallel approach to the KLT implementation is presented. However, in [89] only hyperspectral images with a limited number of spectral bands (up to 8) were considered. This section discusses novel lossy and lossless spectral decorrelation modules for a hyperspectral image compression system on board satellites based on the KLT and Integer KLT transforms, respectively.

The proposed designs are targeted at the SmartFusion System-on-a-Chip (SoC) platform, which incorporates a Flash-based FPGA fabric and a 32-bit hardwired microcontroller. This heterogeneous SoC embedded platform supports a software-hardware co-design approach, which allows tackling the extreme computational complexity of the algorithms via splitting the functionality between a dedicated hardware accelerator and a powerful RISC microprocessor. In particular, the A2F200 SmartFusion SoC has been used in this work, which includes a 32-bit ARM Cortex M-3 Microcontroller Subsystem (MSS). The FPGA logic, which is based on the ProASIC3 device, runs at 50 MHz, while the Cortex M-3 runs at 100 MHz.

The KLT algorithm has been mapped onto the SmartFusionSoC dividing the constituent computational processes between the embedded Cortex M-3 processor and the hardware accelerator on the FPGA fabric using two different approaches as detailed in [87]. A hardware accelerator (co-processor) is built within the FPGA fabric. The frequently occurring operations are performed in the FPGA fabric to accelerate the execution; while the less frequently occurring ones, such as high-level management and tasks scheduling are executed by the embedded Cortex M-3 processor. The BandMean process requires only sequential addition and division operations on a very large set of data and, if implemented on the hardware co-processor, will lead to an intensive exchange of data between the Cortex M-3 and the FPGA fabric, which will consume a significant time. Therefore, they cannot be efficiently implemented on the hardware co-processor. The same applies to the MeanSub process, where only subtraction operations are involved. In the first of the two aforementioned approaches, the co-processor executes the covariance matrix

calculation, the most computationally intensive parts of the Eigenvectors calculation and the multiplication Eigenvectors \times SubMean. In the second approach, only the covariance matrix and the matrix multiplication Eigenvectors \times SubMean are executed within the hardware co-processor. The rationale behind that is to reduce the bit-width of the datapath. In the first approach, the hardware accelerator utilizes a 32-bit data-width. However, while the eigenvectors calculation requires 32-bit operations, both the covariance and the matrix multiplication (Eigen \times SubMean) are performed on a 12-bit wide datapath. Therefore, excluding the eigenvectors computations from the hardware co-processor makes all the operations performed on a 12-bit wide hardware datapath, freeing hardware resources.

The Integer KLT algorithm has been mapped onto the SmartFusionSoC following a similar software-hardware co-design strategy as with the KLT design above. It is found that the computations of the covariance matrix and lifting scheme take the most of the execution time of the Integer KLT. Therefore, by implementing these two processes into the FPGA, a significant acceleration can be achieved. Details of the design can be found in [90].

The KLT and Integer KLT designs, discussed above, lend themselves very well to a combined hardware implementation, in which the designs share some of the computational modules. Such a unified design, performing both functions, would take less hardware resources than the two individual ones. This is possible because the KLT and Integer KLT algorithms are never used simultaneously, due to the different types of the inter-band compression that they realize.

A possible mapping with regards to the SmartFusion SoC is to execute the computations of the covariance matrix and all the matrix multiplications in the FPGA fabric, while the rest of the operations are performed by the Cortex M-3 processor. According to a preliminary estimate of the hardware resources such a combined joint design will only require 10% more system gates and 25% more SRAM blocks than the individual designs.

C. Implementation results

The AVIRIS hyperspectral images are used as test images in the KLT implementation process. In particular, a portion of the AVIRIS Cuprite image¹⁰, comprised of 512 \times 512 pixels, was employed in the experimental work, presented in this section. The KLT design uses a signed 32-bit fixed-point data format, where the first bit represents the sign, the next 16 bits - the integer number and the remaining 15 bits - the fractional part. An error will be accumulated during the processing, the size of which depends on the matrix size and range of the input data. The effect of using this data format on the compression performance compared with a floating point format is an average compression ratio reduction of less than 5% [87].

In order to assess the performance of the proposed system in terms of execution time, the KLT algorithm was implemented on a desktop using an Intel Dual core (2.14 GHz) processor and on the embedded Cortex M-3 processor. The latter was compared

¹⁰ <http://aviris.jpl.nasa.gov/freedata>

with the latency of the proposed SoC system, operating at 100 MHz. The execution times of each process of the KLT algorithm are presented in Table II [85]. The latencies of the PC and Cortex M-3 implementations illustrate the computation intensity of each process, confirming that the matrix multiplication (Eigen \times SubMean) and the covariance matrix computation are the most computationally intensive operations, which together consume more than 95% of the overall processing time.

As it can be seen in Table II, the iterative nature of the eigenvectors computations makes the hardware acceleration very efficient in the first approach, outlined above. In fact, the execution time is reduced by more than 50% while efficient accelerations of 31.6% and 34.3% are achieved on the covariance and the matrix multiplication (Eigen \times SubMean) processes, respectively. In total, the first approach offers a higher execution speed for the overall KLT algorithm by more than 33% when mapped onto the targeted FPGA-based SoC.

The higher level of parallelism offered by the second approach leads to a noticeable further improvement in the performance. As it can be noticed from Table II, an acceleration of 46.3% in the covariance process and 59.4% in the matrix multiplication (Eigen \times SubMean) is achieved, leading to an overall acceleration of 54.3%, cutting the processing time by more than half.

The power consumption of both approaches, estimated using the Actel Smart Power tool, is outlined in Table III, which shows that it is less than 0.25 W, with the first design approach being less power hungry. For the purpose of comparison, the power consumption of the SoC design introduced in [89] was estimated, and it was found to be more than 2W [85], which is in stark contrast to the results in Table III. It can be concluded that the power consumption, offered by the proposed system, is significantly lower despite the four times higher number of processed spectral bands. Table IV summarizes the required hardware resources for the implementation of the KLT co-processor using both proposed approaches (including the AMBA Bus interface). Since Flash FPGA devices usually offer much smaller hardware resources compared to SRAM FPGAs, in both approaches the designs fitted in the FPGA fabric utilizing less than the full amount of gates.

Table V [90] shows the acceleration offered by the proposed Integer KLT design. As it can be seen from this table, the proposed SoC design accomplishes an acceleration of 44.5% compared to the embedded Cortex M3 processor, almost halving the processing time. The overall power consumption is estimated to be less than 0.25 W. The hardware resources required by the Integer KLT design are identical with the resources incurred by the KLT design based on the second approach, as shown in Table IV.

In conclusion, the experimental results presented in this section show that the developed FPGA-based KLT and Integer KLT implementations could be used for spectral decorrelation in hyperspectral compression systems on board remote sensing satellites for certain operational scenarios. However, further research efforts are necessary on reducing the order of complexity of some crucial computations to enable real-time execution of the algorithms.

V. SECOND CASE STUDY: SPECTRAL UNMIXING OF HYPERSPECTRAL IMAGES

In this section, we present spectral unmixing as a hyperspectral image processing case study. More concretely, the implemented unmixing chain is graphically illustrated by a flowchart in Fig. 5 and consists of two main steps: (1) endmember extraction, implemented in this work using the N-FINDR algorithm [56], and (2) non-negative abundance estimation, implemented in this work using ISRA, a technique for solving linear inverse problems with positive constraints [57]. It is worth to mention that we have selected the ISRA algorithm because it provides the following important features: (i) its iterative unmixing nature allows controlling the quality of the obtained solutions depending on the number of iterations executed; (ii) it guarantees convergence after a certain number of iterations; and (iii) it provides positive values in the results of the abundances, which is an important consideration in unmixing applications, as the derivation of negative abundances (which it is possible if an unconstrained model for abundance estimation is applied) is not physically meaningful.

The section is organized as follows. Subsection V.A describes in detail the N-FINDR and the image space reconstruction algorithm (ISRA) used for endmember extraction and abundance estimation purposes, respectively, while Subsection V.B introduces a dynamically reconfigurable FPGA implementation of the considered chain and discloses the most significant results achieved.

A. N-FINDR and the ISRA algorithms

1) The N-FINDR endmember extraction algorithm

In the following, we provide a detailed step-by-step algorithmic description of the original N-FINDR algorithm developed by Winter. It is interesting to notice that the algorithm below represents our own effort to delineate the steps implemented by N-FINDR using available references in the literature [56], [91]. However, it is also worth noting that the N-FINDR algorithm has never been fully disclosed. As a result, this description was developed based on the limited published results available and our own interpretation. Nevertheless, the algorithm below has been verified using the N-FINDR software, provided by the authors, where we have experimentally tested that the software produces essentially the same results as the code below, provided that initial endmembers are generated randomly. The original N-FINDR algorithm can be summarized by the following steps:

a) Feature reduction. Since in hyperspectral data typically the number of spectral bands (n) is much larger than the number of endmembers (p), i.e., $n \gg p$, a transformation that reduces the dimensionality of the input data is required. Hence the first step consist on applying a dimensionality reduction transformation such as the minimum noise fraction [92] or PCA [93] to reduce the dimensionality of the data from n to $p - 1$, where p is an input parameter to the algorithm (number of endmembers to be extracted).

b) Initialization. Let $\{E_1^{(0)}, E_2^{(0)}, \dots, E_p^{(0)}\}$ be a set of endmembers randomly extracted from the input data.

c) Volume calculation. At iteration $k \geq 0$, calculate the volume defined by the current set of endmembers as follows:

$$V(E_1^{(k)}, E_2^{(k)}, \dots, E_p^{(k)}) = \frac{\left| \det \begin{bmatrix} 1 & 1 & \dots & 1 \\ E_1^{(k)} & E_2^{(k)} & \dots & E_p^{(k)} \end{bmatrix} \right|}{(p-1)!} \quad (2)$$

d) Replacement. For each pixel vector $X(i,j)$ in the input hyperspectral data, recalculate the volume by testing the pixel in all p endmember positions, i.e., first calculate $V(X(i,j), E_2(k), \dots, E_p(k))$, then $V(E_1(k), X(i,j), \dots, E_p(k))$, and so on, until $V(E_1(k), E_2(k), \dots, X(i,j))$. If none of the p recalculated volumes is greater than $V(E_1(k), E_2(k), \dots, E_p(k))$, then no endmember is replaced. Otherwise, the combination with maximum volume is retained. Let us assume that the endmember absent in the combination resulting in the maximum volume is denoted by $E_j^{(k+1)}$. In this case, a new set of endmembers is produced by letting $E_j^{(k+1)} = X(i,j)$ and $E_i^{(k+1)} = E_i^{(k)}$ for $i \neq j$. The replacement step is repeated in an iterative fashion, using as much iterations as needed until there are no more replacements of endmembers.

2) The ISRA abundance maps estimation algorithm

Once a set of p endmembers $E = \{e_j\}_{j=1}^p$ has been identified, a positively constrained abundance estimation, i.e. $\Phi_j \geq 0$, can be obtained using ISRA, a multiplicative algorithm based on the following iterative expression:

$$\hat{\Phi}^{k+1} = \hat{\Phi}^k \left(\frac{E^T \cdot X}{E^T E \cdot \hat{\Phi}^k} \right) \quad (3)$$

where the endmember abundances at pixel x are iteratively estimated, so that the abundances at the $k+1$ -th iteration, Φ^{k+1} , depend on the abundances estimated at the k -th iteration, Φ^k . The procedure starts with an unconstrained abundance estimation which is progressively refined in a given number of iterations. For illustrative purposes, Fig. 7 shows the ISRA pseudocode for unmixing one hyperspectral pixel vector x using a set of E endmembers. For simplicity, in the pseudocode x is treated as an n -dimensional vector, and E is treated as an $n \times p$ -dimensional matrix. The estimated abundance vector Φ is a p -dimensional vector, and variable *iters* denotes the number of iterations per pixel in the abundance estimation process. The pseudocode is subdivided into the numerator and denominator calculations in (3). When these terms are obtained, they are divided and multiplied by the previous abundance vector. It is important to emphasize that the calculations of the fractional abundances for each pixel are independent, so they can be calculated simultaneously without data dependencies, thus increasing the possibility of parallelization.

B. FPGA-based linear unmixing of hyperspectral images

With reconfigurable hardware, it is possible to apply much of the flexibility that was formally restricted to software developments only. The idea is that FPGAs can be reconfigured on the fly. This approach is called temporal partitioning [94], [95] or run-time

reconfiguration [96]. Basically, the FPGA (or a region of the FPGA) executes a series of tasks one after another by reconfiguring itself between tasks [97]. The reconfiguration process updates the functionality implemented in the FPGA, and a new task can then be executed. This time-multiplexing approach allows for the reduction of hardware components on-board since one single reconfigurable module can substitute several hardware peripherals carrying out different functions during different phases of the mission.

To implement the spectral unmixing chain on FPGAs, we will take advantage of their reconfigurability. The idea is to start with the implementation of the N-FINDR algorithm occupying the entire FPGA in order to maximize the parallelization of the endmember extraction algorithm. Once the endmembers are found, we will apply dynamic reconfiguration to replace the N-FINDR algorithm with the parallelized ISRA, occupying again the entire FPGA. In our particular case, the N-FINDR and ISRA algorithms are easily scalable because their basic processing units require few hardware resources. Hence, it will be shown that the FPGA is almost full in both cases. More details about our implementations of the N-FINDR and ISRA modules can be found in [55] and [58], although this is the first time that we have used these two modules to build an spectral unmixing chain using run-time reconfiguration. To this end we have included support for inter-task communication following a shared memory scheme. Thus, the N-FINDR module will store its results in a known memory address that correspond to an external DDR2 SRAM memory. Once the system has carried out the reconfiguration, the ISRA module will load the data from that memory. Using external memories to store and load the data may introduce significant delays in the execution. As we will explain in detail, we have overcome this problem including a direct memory access module (DMA) in the system and overlapping the memory transfers with useful computations. With this approach both the reconfiguration and the communications introduce a negligible overhead.

1) N-FINDR hardware design

Fig. 8 describes the general architecture of the hardware used to implement the N-FINDR algorithm, along with the I/O communications. For data input, we use a DDR2 SDRAM and a DMA (controlled by a PowerPC using a prefetching approach) with a FIFO to store pixel data. For data output, we use again the PowerPC to write the position of the endmembers in the DDR2 SDRAM. Finally, the N-FINDR module is used to implement our version of the N-FINDR algorithm. At this point, it is worth to mention that currently we have not developed a hardware implementation for the PCA or the MNF algorithm in order to carry out the dimensionality step. Hence, we will assume that this step has been previously performed and our hardware implementation will carry out all the remaining steps.

The most time consuming part of the algorithm is the volume calculation. The limited available resources in a small or medium FPGA to calculate determinants of large order, makes it difficult to develop an efficient implementation of the algorithm. In our implementation we have solved this issue by taking advantage of the fundamental properties of the determinants and applying

them systematically to transform the determinant in others who are increasingly easy to calculate, down to one that is trivial. For the design of the algorithm we use the matrix triangulation method.

Fig. 9 shows the hardware architecture used to implement the volume calculation step. We use registers to store the pixel vectors selected as endmembers until the current moment, their positions in the image and their volume, and also the current pixel vector data, its position, its greater volume and the index inside the matrix where it is obtained. Moreover, we have included a module that calculates the absolute value of the determinant using the matrix triangulation process: first, for $j = 2, \dots, n$ we take a multiple a_{j1}/a_{11} of the first row and subtract it to the j -th row, to make $a_{j1} = 0$. Thus, we remove all elements of matrix A below the ‘pivot’ element a_{11} in the first column. Now, for $j = 3, \dots, n$, we take a multiple a_{j2}/a_{22} of the second row and subtract it to the j -th row. When we have finished this, all sub-diagonal elements in the second column are zero, and we are ready to process the third column. Applying this process to columns $i = 1, \dots, n-1$ completes the matrix triangulation process and matrix A has been reduced to upper triangular form. These operations are carried out by the data path (see Fig. 10).

Obviously, if one of the diagonal pivots a_{ii} is zero, we cannot use a_{ii} to remove the elements below it; we cannot change a_{ji} by subtracting any multiple of $a_{ii} = 0$ to it. We must switch row i with another row k below it, which contains a nonzero element a_{ki} in the i -th column. Now the new pivot a_{ii} is not zero, and we can continue the matrix triangulation process. If $a_{ki} = 0$ for $k = i, \dots, n$, then it will not be satisfactory to switch row i with any of rows below it, as all the potential pivots are zero and therefore $\det A = 0$. This behavior has been implemented using a modified circular queue with a small control unit (see Fig. 10). Finally, the segmented multiplier calculates the multiplication of the main diagonal elements of the triangular matrix and obtains the absolute value. Table VI summarizes the computational requirements for the N-FINDR hardware implementation, where M and L represent the spatial coordinates, n the number of bands of the hyperspectral image and p the number of endmembers to be extracted.

2) ISRA hardware design

Fig. 11 describes the general architecture of the hardware used to implement the ISRA, along with the I/O communications, following a similar scheme to the previous subsection. For data input, we use a DDR2 SDRAM and a DMA (controlled by a PowerPC using a prefetching approach) with a FIFO to store pixel data. ISRA module is used to implement our parallel version of the ISRA. Finally, a transmitter is used to send the fractional abundances via a RS232 port.

Fig. 12 shows the hardware architecture used to implement the ISRA module. Three different memories are used to store the endmembers, the current pixel, and the fractional abundances for the current pixel, respectively. The ISRA data path represents the hardware architecture used to perform the calculations. The control unit carries out the ISRA execution: it reads the appropriate memory locations for each of the memories and updates the fractional abundances. In addition, we use a combinational circuit based on multiplexers to select the appropriate input data. Once calculated, the system writes the estimated

abundances to the read FIFO.

Fig. 12 also describes the architecture of the data path used to implement the ISRA module. The dot-product unit is used for calculating both the numerator and denominator in Eq. (4), allowing a proper reuse of hardware resources. To perform the update of a fractional abundance value, it proceeds as follows: during n cycles (where n is the number of bands) it computes the dot-product between the current pixel and the endmember corresponding to the proportion of abundance that is being updated. In the next clock cycle, the result of the dot-product is multiplied by the previous abundance fraction and the result is stored in register num , thus concluding the calculation of the numerator. To calculate the denominator, the aforementioned procedure is repeated p times (where p is the number of endmembers) with the appropriate input data, while partial results are accumulated using an adder and the register den . The calculation of the denominator requires therefore $p \times (n+1)$ clock cycles. The process finalizes with the division between the numerator and the denominator in 5 clock cycles. The computational requirements of this design are presented in Table VI, where i represents the number of iterations per pixel in the abundance estimation process.

3) Implementation results

The hardware architectures previously described have been implemented on a ML410 board, a reconfigurable board with a single Virtex-4 XC4VFX60 FPGA component, a memory slot which holds up to 2GBytes and some additional components not used in our implementation. We have used a Xilinx Virtex-4 XC4VFX60 FPGA because its features are very similar to the space-grade Virtex-4QV XQR4VFX60 FPGA. As far as the results obtained by our FPGA-based unmixing chain were exactly the same than the ones obtained with its equivalent software versions, and the efficiency in terms of unmixing of the N-FINDR and the ISRA algorithms have been already proven in many state-of-the-art works, we have decided to center this subsection in the hardware implementation results achieved.

Table VII shows the resources used for our hardware implementation of the proposed N-FINDR algorithm design optimized to extract up to 21 endmembers, and Table VIII shows the resources used for our hardware implementation of the proposed ISRA design for 16 ISRA basic units in parallel, conducted on the Virtex-4 XC4VFX60 FPGA of the ML410 board. This FPGA has a total of 25280 slices, 50560 slice flip flops, and 50560 four input look-up tables available. In addition, the FPGA includes some heterogeneous resources, such as two PowerPCs, 128 DSP48Es, and distributed Block RAMs. Table IX outlines the power consumption of both designs, estimated using the Xilinx Power Estimator tool. Taking into account that we are using most of the FPGA computing resources simultaneously, these numbers are very affordable, especially when compared to the power numbers of high-performance general-purpose processors or GPUs. Finally, Table X reports the processing times measured for two well-known hyperspectral data sets not only for our FPGA implementation, but also for an equivalent software version developed in C language and executed on a PC with an Intel Core i7 processor at 2.2 Gigahertz and 4 Gigabytes of RAM. The first dataset corresponds with a portion of 350×350 pixels of the AVIRIS Cuprite scene also used for the compression case study reported in

the previous section of this paper, while the second data set corresponds to an EO-1 Hyperion data set available in radiance units collected over the same Cuprite mining district as the aforementioned AVIRIS scene. In this case, we used a full EO-1 Hyperion flightline with much larger dimensions, i.e., 6479×256 pixels and 242 spectral bands, and a total size of around 800 Megabytes. In both cases our hardware implementation achieves a speedup of 9x which is a remarkable result if it is taken into account that the Intel Core i7 uses a much more recent technology (our Virtex-4 uses 2004 technology whereas the Intel Core i7 uses technology from the end of 2011), and is a power hungry platform that demands complex heat dissipation systems, which are not needed for our FPGA design.

To conclude this section, we emphasize that our reported FPGA processing times are still slower than the sensors that collect the hyperspectral data. For instance, the cross-track line scan time in EO-1 Hyperion, a push-broom instrument, is quite fast (4.5 ms to collect 256 full pixel vectors). This introduces the need to process a hyperspectral data set with 256×6479 pixels (such as our considered EO-1 Hyperion Cuprite scene) in 29.16 s to fully achieve real-time performance. Although we believe that the inclusion of more recent FPGA boards could significantly improve our reported processing times, we have decided to report results on a board which is very close to those certified for space operation in the hope of illustrating the real challenges to be faced when porting hyperspectral imaging algorithms into certified FPGA hardware. For example, in a Virtex-4 FPGA XQR4VLX200 (89088 slices) certified for space, we could have 3.5 times more ISRA units in parallel simply synthesizing the second stage of the hardware architecture for the new number of units in parallel. In this way, we could achieve a speedup of 3.5x with regards to our current implementation without any modifications in our proposed design. In the case of an airborne platform without the need for space-certified hardware, we could use a Virtex-6 XQ6VLX550T (550000 logic cells) which has nearly 10 times more logic cells than the FPGA used in our experiments. In future developments, we will also focus on improving our proposed implementation to achieve a better utilization of hardware resources and reduce the reported processing times, which in any event are considered to be acceptable in many remote sensing applications.

VI. SUMMARY AND FUTURE RESEARCH

This paper discussed the role of reconfigurable hardware in remote sensing missions by providing an extensive review and analysis of the (current and future) capabilities of FPGA devices, which are now the standard for on-board processing and analysis of remotely sensed data. Both the technological aspects of FPGA hardware and also specific implementation issues are covered in the context of two specific case studies dealing with remotely sensed hyperspectral data compression and spectral unmixing. While many other techniques currently exist for hyperspectral data exploitation, the selection of these two applications was based on their relevance and fast development in recent years. Several summarizing statements can be formulated after our detailed assessment:

- FPGAs currently represent the best choice for on-board hyperspectral processing due to their compact size and low weight, particularly when compared with other specialized hardware devices such as GPUs, which exhibit much higher power dissipation figures. FPGAs also perform well in high-throughput signal processing tasks encountered in space, like processing and/or compressing a hyperspectral image.
- An important property of FPGAs is their inherent ability to change their functionality- through partial or full reconfiguration. In other words, they permit changes to the usage model and the data processing paradigm in space rather than hard-coding of all components prior to launch. Although FPGAs that support run-time reconfiguration has been available for many years, initially the reconfiguration flow was too complex, and very few designers' teams use it. In fact the example described in Section V is the first time where run-time reconfiguration has been used for hyperspectral images. However, in the recent years FPGA vendors have done an important effort to simplify the development of run-time reconfigurable systems. Hence future on-board processing systems could easily take advantage of this useful feature.
- Another important technological advantage of FPGAs over commodity high performance platforms such as current GPUs is that FPGAs can be manufactured in order to resist high levels of radiation without changing the content of their inner memories, which determines the FPGA programming. Moreover, since FPGAs implement custom designs, additional fault-tolerant levels can be included in the system when needed, such as dual or triple modular redundancy.
- The aforementioned properties make FPGAs a very attractive solution for advanced hyperspectral data exploitation applications such as (lossless and lossy) hyperspectral data compression and unmixing, as discussed in this work. However, FPGAs cannot still provide real-time processing performance on board in the considered analysis scenarios. This is mainly due to the additional hardware resources required by the reconfigurability property, and also to the extra resources needed to enable adaptive execution of a particular processing algorithm among a suite of algorithms, implementing a desired functionality. In addition, in view of the computationally intensive nature of the hyperspectral imaging algorithms new low-complexity mapping hardware methods are required enabling real-time execution based on an appropriate trade-off among design requirements. In this regard, it is important to mention that due to the time needed for the certification process of hardware, current space-certified FPGA devices do not represent the state-of-the-art hardware components, hence the technology available onboard is generally one generation behind the latest one available in the market.
- Another consideration when porting processing algorithms to FPGAs is the programming effort needed, which is generally higher than the one required to develop an efficient implementation in a GPU or multi-core processor since

the full processing architecture needs to be developed from scratch. High-level design tools such as Handel-C¹², Catapult C¹³, Dime-C¹⁴, or Impulse C¹⁵ among others have simplified the design process for FPGA developers, but the learning curve for inexperienced users implementing their algorithms in FPGA technology is still significant as compared to the efforts needed in other kind of commodity high performance computing platforms such as GPUs or multi-cores. In this sense, some efforts have been done in order to directly map hyperspectral imaging algorithms described in MATLAB (which is by far the preferred programming language within the hyperspectral research community) onto a generic FPGA, but the results are still far from being efficient when compared with the classical FPGA design flow [79], [98]. In any case, all the FPGA vendors are trying to reduce the design complexity not only including support to synthesize high-level language such as C or C++, but also including large component intellectual-property (IP) libraries with hundreds of frequently used designs, and allowing the easy integration of previously designed hardware blocks in a custom design.

In summary, reconfigurable hardware offers an unprecedented opportunity to achieve the long-desired goal of being able to perform real-time processing of high-dimensional remote sensing data (such as hyperspectral images) on board instrument. FPGA technology is now highly consolidated in both airborne and spaceborne scenarios, and provides an important property such reconfigurability which allows for a longer exploitation of the hardware and its dynamic adaptation to the availability of new processing techniques and algorithms. Due to these reasons, FPGA technology has attracted the attention of international space agencies for their future on-board processing systems. A good example is the SpaceCube program of the NASA Goddard Space Flight Center, which aims to provide 10x to 100x improvements in on-board computing power while lowering relative power consumption and cost thanks to the incorporation of commercially available radiation-tolerant FPGAs to their flight processing systems [99].

In this regard, the review provided in this contribution analyzes current and future developments in this area and further discusses the role of FPGAs in future spaceborne missions for Earth observation, intended to provide full coverage of our planet with extremely high spatial, spectral and temporal resolutions. Expected developments in this field will comprise the full integration of reconfigurable technology into spaceborne missions and a more dynamic certification process by international agencies that will allow for the integration of state-of-the-art hardware devices into remote sensing missions as soon as they become available. Further developments towards enhancing the programmability of FPGA devices are also expected in order to facilitate the porting of algorithms to this kind of hardware using high-level tools.

¹² <http://www.handelc.com>

¹³ http://www.calypto.com/catapult_c_synthesis.php

¹⁴ <http://www.nallatech.com/Development-Tools/dime-c.html>

¹⁵ <http://www.impulseaccelerated.com/>

ACKNOWLEDGMENTS

The authors would like to acknowledge the Guest Editors of this special issue for their very kind invitation to provide a contribution, as well as the three anonymous reviewers for their outstanding comments and suggestions for improvement. Additionally, funding from the European Commission in the framework of the TOLOMEO (FP7-PEOPLE-2010-IRSES) project and from the Spanish Government in the framework of the CEOS-SPAIN (AYA2011-29334-C02-02), DREAMS (TEC2011-28666-C04-04), AYA2009-13300-C03-02, TIN2009-09806, and TIN2010-21291-C02-01 projects is gratefully acknowledged.

REFERENCES

- [1] R. Trautner, "ESA's Roadmap for Next Generation Payload Data Processors," *DASIA Conference*, 2011. Available: <http://www.esa.int/TEC/OBDP/>
- [2] Zynq-7000 All Programmable SoC. Available: <http://www.xilinx.com/products/silicon-devices/soc/zynq-7000/index.htm>
- [3] Altera SoC FPGA. Available: <http://www.altera.com/devices/processor/soc-fpga/proc-soc-fpga.html>
- [4] S. Hauck and A. Dehon, Eds., "Device architecture," in *Reconfigurable Computing*, Elsevier, 2008, pp. 3-27
- [5] Space-grade Virtex-4QV FPGA. Available: <http://www.xilinx.com/products/silicon-devices/fpga/virtex-4qv/index.htm>
- [6] Space-grade Virtex-5QV FPGA. Available: <http://www.xilinx.com/products/silicon-devices/fpga/virtex-5qv/index.htm>
- [7] Reports on radiation effects in Xilinx FPGAs. Available: <http://parts.jpl.nasa.gov/organization/group-5144/radiation-effects-in-fpgas/xilinx/>
- [8] ProASIC3 FPGA. Available: <http://www.actel.com/products/pa3/>
- [9] T. Kuwahara, "FPGA-based Reconfigurable On-board Computing Systems for Space Applications," PhD Dissertation, Faculty of Aerospace Engineering and Geodesy of the University of Stuttgart, 2009.
- [10] A. Plaza, Q. Du, Y.-L. Chang and R. L. King, "High Performance Computing for Hyperspectral Remote Sensing", *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 4, no. 3, pp. 528-544, September 2011.
- [11] A. Plaza and C.-I. Chang, *High Performance Computing in Remote Sensing*, CRC Press, 2007.
- [12] A. Plaza, D. Valencia and J. Plaza, "High-Performance Computing in Remotely Sensed Hyperspectral Imaging: The Pixel Purity Index Algorithm as a Case Study," *International Parallel and Distributed Processing Symposium (IPDPS)*, 2006.
- [13] A. Plaza, J. Plaza, A. Paz and S. Sanchez, "Parallel Hyperspectral Image and Signal Processing," *IEEE Signal Processing Magazine*, vol. 28, no. 3, pp. 119-126, May 2011.
- [14] C. A. Lee, S. D. Gasster, A. Plaza, C.-I. Chang, B. Huang, "Recent Developments in High Performance Computing for Remote Sensing: A Review", *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 4, no. 3, pp. 508-527, September 2011.

- [15] G. Yu, T. Vladimirova and M.N. Sweeting, "Image Compression Systems on Board Satellites," *Acta Astronautica*, vol. 64, no. 9-10, pp. 988-1005, May 2009.
- [16] N.S. Jayant and Peter Noll, *Digital Coding of Waveforms: Principles and Applications to Speech and Video*, Prentice-Hall Inc., 1984.
- [17] M. Rabbani and P. W. Jones, *Digital Image Compression Techniques*, SPIE Press, 1991.
- [18] B.V. Brower, D. Couwenhoven, B. Gandhi, C. Smith, "ADPCM for advanced LANDSAT downlink applications," *Conference Record of The Twenty-Seventh Asilomar Conference on Signals, Systems and Computers*, vol.2, pp.1338-1341, 1993.
- [19] CCSDS (1997). "Lossless Data Compression, Recommendation for space data system standards," CCSDS 121.0-B-1, May 1997
- [20] W. B. Pennebaker and J. L. Mitchell, *JPEG Still Image Data Compression Standard*, Chapman & Hall, 1993.
- [21] M. J. Weinberger, G. Seroussi, and G. Sapiro, "The LOCO-I Lossless Image Compression Algorithm: Principles and Standardization into JPEG-LS," *IEEE Transactions on Image Processing*, vol. 9, no. 8, pp. 1309–1324, August 2000.
- [22] D. S. Taubman, M.W. Marcellin, *JPEG2000: Image Compression Fundamentals, Standard and Practice*, Kluwer, 2002.
- [23] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Transactions on Signal Processing*, vol. 41, pp. 3445–3462, Dec. 1993.
- [24] A. Said and W. A. Pearlman, "A new fast and efficient image codec based on set partitioning in hierarchical trees," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, pp. 243–250, June 1996.
- [25] CCSDS (2005). "Image Data Compression, Recommendation for space data system standards", CCSDS 122.0-B-1, November 2005.
- [26] P.-S. Yeh and J. Venbrux, "A high performance image data compression technique for space applications," *NASA Earth Science Technology Conference*, 2003.
- [27] J. Serra-Sagrasta, C. Fernandez-Cordoba, F. Auli-Llinas, F. Garcia-Vilchez, J. Minguillon, "Lossy coding techniques for high resolution images," *Proceedings of SPIE*, vol. 5238, pp. 276-287, 2004.
- [28] J. Serra-Sagrasta, Joan, F. Auli-Llinas, F. Garcia-Vilchez, C. Fernandez-Cordoba, "Review of CCSDS-ILDC and JPEG2000 coding techniques for remote sensing," *Proceedings of SPIE*, vol.5573, pp. 250-261, 2004.
- [29] Y. Pen-Shu, P. Armbruster, A. Kiely, B. Masschelein, G. Moury, C. Schaefer, C. Thiebaut, "The New CCSDS Image Compression Recommendation", *IEEE Conference on Aerospace*, pp. 1- 8, 2005.
- [30] N.R. Mat Noor and T. Vladimirova, "Investigation into Lossless Hyperspectral Image Compression for Satellite Remote Sensing," to appear in *International Journal of Remote Sensing*.

- [31] G. Motta, F. Rizzo, and J. A. Storer, Eds., "Lossless Predictive Compression of Hyperspectral Images," in *Hyperspectral Data Compression*, Springer-Verlag New York, 2005.
- [32] M. J. Ryan, J.F. Arnold, "The Lossless Compression of AVIRIS Images by Vector Quantization," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 35, no. 3, pp. 546-550, May 1997.
- [33] M. J. Ryan, J.F. Arnold, "Lossy Compression of Hyperspectral Data Using Vector Quantization," *Remote Sensing of Environment*, vol. 61, no. 3, pp. 419-436, September 1997.
- [34] G. Motta, F. Rizzo, and J. A. Storer, Eds., "An Architecture for the Compression of Hyperspectral Imagery," in *Hyperspectral Data Compression*, Springer-Verlag New York, 1995.
- [35] J. A. Saghri, A. G. Tescher, J. T. Reagan, "Practical Transform Coding of Multispectral Imagery," *IEEE Signal Processing Magazine*, vol. 12, no. 1, pp. 32-43, January 1995.
- [36] G. Liu, F. Zhao, "Efficient Compression Algorithm for Hyperspectral Images Based on Correlation Coefficients Adaptive 3D Zerotree Coding," *IET Image Processing*, vol. 2, no. 2, pp. 72-82, April 2008.
- [37] I. Blanes, J. Serra-Sagrsta, "Clustered Reversible-KLT for Progressive Lossy-to-Lossless 3D Image Coding," *Data Compression Conference(DCC)*, pp. 233-242, 2009.
- [38] CCSDS (2012b). "Lossless Multispectral & Hyperspectral Image Compression - Recommended Standard," CCSDS 123.0-B-1, May 2012.
- [39] CCSDS (2012a). "Lossless Data Compression - Recommended Standard," CCSDS 121.0-B-2, May 2012.
- [40] Y.-T. Hwang, C.-C. Lin, R.-T. Hung, "Lossless Hyperspectral Image Compression System-Based on HW/SW Codesign," *IEEE Embedded Systems Letters*, vol.3, no.1, pp.20-23, March 2011.
- [41] N. Aranki, D. Keymeulen, A. Bakhshi, M. Klimesh, "Hardware Implementation of Lossless Adaptive and Scalable Hyperspectral Data Compression for Space," *NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, pp.315-322, 2009.
- [42] G. Yu, T. Vladimirova, M. N. Sweeting, "FPGA-based on-board multi/hyperspectral image compression system," *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 5, pp.212-215, 2009.
- [43] A. C. Miguel, A. R. Askew, A. Chang, S. Hauck, R. E. Ladner, E. A. Riskin, "Reduced complexity wavelet-based predictive coding of hyperspectral images for FPGA implementation," *Data Compression Conference*, pp. 469- 478, 2004.
- [44] D. Valencia, A. Plaza, "FPGA-Based Hyperspectral Data Compression Using Spectral Unmixing and the Pixel Purity Index Algorithm", in *Proceedings of the 6th International Conference on Computational Science(ICCS)*, pp. 888-891, 2006.
- [45] A. Said, W. A. Pearlman, "A New, Fast, and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees," *IEEE Transactions on Circuits and Systems*, vol. 6, no. 3, pp. 243-250, June 1996.

- [46] A. Plaza, "Towards Real-Time Compression of Hyperspectral Images Using Virtex-II FPGAs," *International Euro-Par Conference*, pp. 248-257, 2007.
- [47] A. Plaza, S. Sanchez, A. Paz and J. Plaza, "GPUs versus FPGAs for Onboard Compression of Hyperspectral Data," *International Workshop on On-Board Payload Data Compression (OBPDC)*, 2010.
- [48] S. Deo, "Power Consumption Calculation of AP-DCD Algorithm Using FPGA Platform," *International Conference on Reconfigurable Computing and FPGAs (ReConFig)*, pp.388-393, 2010.
- [49] A. Plaza, J. A. Benediktsson, J. Boardman, J. Brazile, L. Bruzzone, G. Camps-Valls, J. Chanussot, M. Fauvel, P. Gamba, J. A. Gualtieri, M. Marconcini, J. C. Tilton, G. Trianni, "Recent advances in techniques for hyperspectral image processing," *Remote Sensing of Environment*, vol. 113, supplement 1, pp. 110–122, September 2009.
- [50] J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader and J. Chanussot, "Hyperspectral Unmixing Overview: Geometrical, Statistical and Sparse Regression-Based Approaches," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 2, pp. 354-379, April 2012.
- [51] A. Plaza, P. Martinez, R. Perez, J. Plaza, "A quantitative and comparative analysis of endmember extraction algorithms from hyperspectral data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 3, pp. 650–663, March 2004.
- [52] J. M. P. Nascimento, J. M. Bioucas-Dias, "Does independent component analysis play a role in unmixing hyperspectral data?," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, no. 1, pp. 175–187, January 2005.
- [53] D. Heinz, C.-I. Chang, "Fully constrained least squares linear mixture analysis for material quantification in hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 3, pp. 529–545, March 2001.
- [54] M. Velez-Reyes, A. Puetz, M. P. Hoke, R. B. Lockwood, S. Rosario, "Iterative algorithms for unmixing of hyperspectral imagery," in *Proceedings of SPIE*, vol. 5093, pp. 418–429, 2003.
- [55] C. Gonzalez, D. Mozos, J. Resano and A. Plaza, "FPGA Implementation of the N-FINDR Algorithm for Remotely Sensed Hyperspectral Image Analysis," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, no. 2, pp. 374-388, February 2012.
- [56] M. E. Winter, "N-FINDR: An algorithm for fast autonomous spectral end-member determination in hyperspectral data," *Proceedings of SPIE*, vol. 3753, pp. 266-275, 1999.
- [57] S. Torres-Rosario, "Iterative algorithms for abundance estimation on unmixing of hyperspectral imagery," Master Thesis, University of Puerto Rico, 2004.
- [58] C. Gonzalez, J. Resano, A. Plaza and D. Mozos, "FPGA Implementation of Abundance Estimation for Spectral Unmixing of Hyperspectral Data Using the Image Space Reconstruction Algorithm", *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 1, pp. 248-261, February 2012.

- [59] C. Gonzalez, D. Mozos, J. Resano and A. Plaza, "FPGA for Computing the Pixel Purity Index Algorithm on Hyperspectral Images," *Engineering of Reconfigurable Systems and Algorithms (ERSA)*, 2010.
- [60] C. Gonzalez, J. Resano, D. Mozos, A. Plaza and D. Valencia, "FPGA Implementation of the Pixel Purity Index Algorithm for Remotely Sensed Hyperspectral Image Analysis," *EURASIP Journal on Advances in Signal Processing*, vol. 2010, Article ID 969806, 13 pages, 2010.
- [61] J. Boardman, "Automating spectral unmixing of AVIRIS data using convex geometry concepts," *Summaries of Airborne Earth Science Workshop*, JPL Publication 93-26, pp. 111-114, 1993.
- [62] D. Valencia, A. Plaza, M. A. Vega-Rodríguez, and R. M. Pérez, "FPGA design and implementation of a fast pixel purity index algorithm for endmember extraction in hyperspectral imagery," *Chemical and Biological Standoff Detection III, Proceedings of SPIE*, 2005.
- [63] J. Morales, N. Medero, N. G. Santiago, J. Sosa, "Hardware Implementation of Image Space Reconstruction Algorithm using FPGAs," *49th IEEE International Midwest Symposium on Circuits and Systems (MWSCAS)*, vol.1, pp.433-436, 2006.
- [64] C.-I. Chang, W. Xiong, C.-C. Wu, "Field-Programmable Gate Array Design of Implementing Simplex Growing Algorithm for Hyperspectral Endmember Extraction," *IEEE Transactions on Geoscience and Remote Sensing*. Accepted for inclusion in a future issue of this journal.
- [65] C.-I Chang, C. Wu, W. Liu, Y. C. Ouyang, "A growing method for simplex-based endmember extraction algorithms," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 10, pp. 2804-2819, October 2006.
- [66] C.-I Chang, C. C. Wu, C.-S. Lo, M.-L. Chang, "Real-time simplex growing algorithms for hyperspectral endmember extraction," *IEEE Transactions on Geoscience and Remote Sensing*., vol. 40, no. 4, pp. 1834-1850, April 2010.
- [67] W. Xiong, C.-C. Wu, C.-I Chang, K. Kapalkis, H. M. Chen, "Fast algorithms to implement N-FINDR for hyperspectral endmember extraction," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 4, no. 3, pp. 545-564, September 2011.
- [68] S. Lopez, P. Horstrand, G.M. Callico, J.F. Lopez and R. Sarmiento, "A Novel Architecture for Hyperspectral Endmember Extraction by Means of the Modified Vertex Component Analysis (MVCA) Algorithm," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*. Accepted for inclusion in a future issue of this journal.
- [69] S. Lopez, P. Horstrand, G.M. Callico, J.F. Lopez and R. Sarmiento, "A Low-Computational-Complexity Algorithm for Hyperspectral Endmember Extraction: Modified Vertex Component Analysis," *IEEE Geoscience and Remote Sensing Letters*, vol. 9, no. 3, pp. 502-506, May 2012.
- [70] D. Hongtao, Q. Hairong, "An FPGA implementation of parallel ICA for dimensionality reduction in hyperspectral images," *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, vol.5, pp. 3257- 3260, 2004.

- [71] M. Lennon, G. Mercier, M. C. Mouchot and L. Hubert-Moy, "Independent Component Analysis as a tool for the dimensionality reduction and the representation of hyperspectral images," *SPIE Remote Sensing*, vol.4541, pp.2893-2895, 2001.
- [72] A.B. Lim, J.C. Rajapakse and A.R. Omondi, "Comparative Study of Implementing ICNNs on FPGAs," *International Joint Conference on Neural Networks*, vol.1, pp.177-182, 2001.
- [73] F. Sattar and C. Charayaphan, "Low-cost design and implementation of an ICA-based blind source separation algorithm," *Annual IEEE International ASIC/SOC Conference*, pp.15-19, 2002.
- [74] A. Jacobs, C. Conger, A. D. George, "Multiparadigm Space Processing for Hyperspectral Imaging," *IEEE Aerospace Conference*, vol. 1, pp.8-11, 2008.
- [75] C. Chang, H. Ren, and S. Chiang, "Real-time processing algorithms for target detection and classification in hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 4, pp. 760- 768, April 2001.
- [76] C.-I Chang, *Hyperspectral Imaging: Techniques for Spectral Detection and Classification*. Kluwer Academic/Plenum Publishers, New York, 2003..
- [77] R. Nekovei, M. Ashtijou, "Reconfigurable acceleration for hyperspectral target detection," *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp.3229-3232, 2007.
- [78] Q. Du and H. Ren, "Real-time linear constrained discriminant analysis to hyperspectral imagery," *Proceeding of SPIE International Conference on Multispectral Image Processing and Pattern Recognition*, vol. 4548, pp. 103-108, 2001.
- [79] S. Bernabe, S. Lopez, A. Plaza, R. Sarmiento, P. G. Rodriguez, "FPGA Design of an Automatic Target Generation Process for Hyperspectral Image Analysis," *IEEE International Conference on Parallel and Distributed Systems (ICPADS)*, pp. 1010-1015, 2011.
- [80] K.-S. Park, S. H. Cho, S. Hong and W.-D. Cho, "Real-time target detection architecture based on reduced complexity hyperspectral processing," *EURASIP Journal on Advances in Signal Processing*, vol. 2008, Article ID 438051, pp. 1-18, September 2008.
- [81] Q. Du and R. Nekovei, "Fast real-time onboard processing of hyperspectral imagery for detection and classification," *Journal of Real-Time Image Processing*, vol. 4, no. 3, pp. 273-286, August 2009.
- [82] Z. K. Baker, M. B. Gokhale, J. L. Tripp, "Matched Filter Computation on FPGA, Cell and GPU," *IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pp.207-218, 2007.
- [83] Q. Du, J.E. Fowler, "Hyperspectral Image Compression Using JPEG2000 and Principal Component Analysis," *IEEE Geoscience and Remote Sensing Letters*, vol. 4,no. 2, pp. 201-205, April 2007.

- [84] N.R. Mat Noor, T. Vladimirova, M.N. Sweeting, "High-Performance Lossless Compression for Hyperspectral Satellite Imagery," *Proc. of 2010 UK Electronics Forum*, pp. 78-83, 2010.
- [85] C. Egho, T. Vladimirova, M. N. Sweeting, "Acceleration of Karhunen-Loève Transform for System-on-Chip Platforms," *Proc. of 7th NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, pp. 272-279, 2012.
- [86] N. R. Mat Noor, T. Vladimirova, "Integer KLT Design Space Exploration for Hyperspectral Satellite Image Compression," *Lecture Notes in Computer Science*, vol. 6935, pp. 661-668, 2011.
- [87] C. Egho, T. Vladimirova, "Hardware Acceleration of Karhunen-Loeve Transform for Compression of Hyperspectral Satellite Imagery," *Australian Space Science Conference (ASSC)*, pp. 237-248, 2011.
- [88] P. Hao, Q. Shi, "Matrix Factorizations for Reversible Integer Mapping", *IEEE Transactions on Signal Processing*, vol. 49, no. 10, pp. 2314-2324, October 2010.
- [89] M. Fleury, R. P. Self, A. C. Downton, "Development of a Fine-Grained Karhunen-Loeve Transform," *Journal of Parallel and Distributed Computing*, vol. 64, no. 4, pp. 520-535, April 2004.
- [90] C. Egho, T. Vladimirova, "Hardware Acceleration of the Integer Karhunen-Loève Transform Algorithm for Satellite Image Compression", *IEEE International Geoscience & Remote Sensing Symposium (IGARSS)*, pp. 4062 – 4065, 2012.
- [91] M. E. Winter, "A proof of the N-FINDR algorithm for the automated detection of endmembers in a hyperspectral image," in *Proceedings of SPIE*, vol. 5425, pp. 31–41, 2004.
- [92] A. A. Green, M. Berman, P. Switzer, M. D. Craig, "A transformation for ordering multispectral data in terms of image quality with implications for noise removal," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 26, no. 1, pp. 65–74, January 1988.
- [93] R. A. Schowengerdt, *Remote Sensing: Models and Methods for Image Processing*, New York: Academic Press, 1997.
- [94] J. Tabero, H. Mecha, J. Septién, S. Román, D. Mozos, "A vertex-list approach to 2D Hw multitasking management in RTR FPGAs," *Proceedings of DCIS conference*, pp. 545–550, 2003.
- [95] S. Román, J. Septién, H. Mecha, D. Mozos, "Constant complexity management of 2D HW multitasking in run-time reconfigurable FPGAs," *Lecture Notes in Computer Science*, vol. 3985, pp. 187–192, 2006.
- [96] J. Resano, J. A. Clemente, C. González, D. Mozos, F. Catthoor, "Efficiently scheduling runtime reconfigurations," *ACM Transactions on Design Automation of Electronic Systems*, vol. 13, no. 4, pp. 58–69, September 2008.
- [97] J. A. Clemente, C. González, J. Resano, D. Mozos, "A task graph execution manager for reconfigurable multi-tasking systems," *Microprocessors and Microsystems*, vol. 34, no. 2–4, pp. 73–83, March 2010.

- [98]S. Lopez, G. M. Callico, A. Medina, J. F. Lopez, R. Sarmiento, “High-Level FPGA-Based Implementation of a Hyperspectral Endmember Extraction Algorithm,” *Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2012.
- [99]T. Flatley, “SpaceCube: A family of reconfigurable hybrid on-board science data processors,” *NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*, Keynote Address I, 2012.



Sebastian Lopez was born in Las Palmas de Gran Canaria, Spain, in 1978. He received the Electronic Engineer degree by the University of La Laguna in 2001, obtaining regional and national awards for his CV during his degree. He got his PhD degree by the University of Las Palmas de Gran Canaria in 2006, where he is actually an Associate Professor, developing his research activities at the Integrated Systems Design Division of the Institute for Applied Microelectronics (IUMA). He is a member of the IEEE Geoscience & Remote Sensing and Consumer Electronics Societies as well as a member of the Publications Review Committee of the IEEE Transactions on Consumer Electronics. Additionally, he currently serves as an active reviewer of the IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS), IEEE Geoscience and Remote Sensing Letters, IEEE Transactions on Circuits and Systems for Video Technology, the Journal of Real Time Image Processing, Microprocessors and Microsystems: Embedded Hardware Design (MICPRO), and the IET Electronics Letters. He has published more than 50 papers in international journals and conferences. His research interests include real-time hyperspectral imaging systems, reconfigurable architectures, video coding standards, and hyperspectral image compression systems.



Tanya Vladimirova, MEng, MSc, PhD, CEng, MIET, MEEE, received an M.Sc. in Applied Mathematics from the Technical University of Sofia, Bulgaria, an MEng in Computer Systems Engineering, and a Ph.D. in VLSI Design from the St. Petersburg Electro-Technical University (LETI), Russia. She is currently Professor of Embedded Systems and VLSI Design at the Department of Engineering, University of Leicester, UK. Prior to that she was with the Department of Electronic Engineering at the University of Surrey and led the VLSI Design and Embedded Systems research group at the Surrey Space Centre. Her research interests are in the areas of low-power on-board integrated circuit design, image processing, intelligent embedded systems and space-based wireless sensor networks. She acted as a co-chair of the NASA conference on Military and Aerospace Applications of Programmable Logic Devices (MAPLD) from 2000 to 2006.



Carlos González received the M.S. and Ph.D. degrees in computer engineering from the Complutense University of Madrid, Madrid, Spain, in 2008 and 2011, respectively. He is currently a Teaching Assistant in the Department of Computer Architecture and Automation of the Universidad Complutense Madrid. As a research member of GHADIR group, he mainly focuses on applying run-time reconfiguration in aerospace applications. His research interests include remotely sensed hyperspectral imaging, signal and image processing, and efficient implementation of large-scale scientific problems on reconfigurable hardware. He is also interested in the acceleration of artificial intelligence algorithms applied to games. Dr. González won the Design Competition of the IEEE International Conference on Field Programmable Technology in 2009 (FPT'09) and in 2010 (FPT'10). He received the Best Paper Award of an Engineer under 35 years old in the International Conference on Space Technology in 2011.



Javier Resano received the Bachelor Degree in Physics in 1997, a Master Degree in Computer Science in 1999, and the Ph.D. degree in 2005 at the Universidad Complutense of Madrid, Spain. Currently he is Associate Professor at the Computer Eng. Department of the Universidad of Zaragoza, and he is a member of the GHADIR research group, from Universidad Complutense, and the GAZ research group, from Universidad de Zaragoza. He is also member of the Engineering Research Institute of Aragon (I3A). His research has been focused in hardware/software co-design, task scheduling techniques, Dynamically Reconfigurable Hardware and FPGA design. His FPGA designs have received several international awards including the first prize in the Design Competition of the IEEE International Conference on Field Programmable Technology in 2009 (FPT'09) and in 2010 (FPT'10).



Daniel Mozos is a permanent professor in the Department of Computer Architecture and Automatics of the Complutense University of Madrid, where he leads the GHADIR research group on dynamically reconfigurable architectures. His research interests include design automation, computer architecture, and reconfigurable computing. Mozos has a B.S. in physics and a Ph.D. in computer science from the Complutense University of Madrid.



Antonio Plaza (M'05–SM'07) received the M.S. and Ph.D. degrees in computer engineering from the University of Extremadura, Spain. He has been a Visiting Researcher with the Remote Sensing Signal and Image Processing Laboratory (RSSIPL), University of Maryland Baltimore County; with the Applied Information Sciences Branch, NASA Goddard Space Flight Center (GSFC), Greenbelt, MD; with the Airborne Visible Infrared Imaging Spectrometer Data Facility, NASA Jet Propulsion Laboratory (JPL), Pasadena, CA; with the Telecommunications and Remote Sensing Laboratory, University of Pavia, Italy; and with the GIPSA-lab, Grenoble Images Parole Signal Automatique, France. He is currently an Associate Professor (with accreditation for Full Professor) with the Department of Technology of Computers and Communications, University of Extremadura, where he is the Head of the Hyperspectral Computing Laboratory (HyperComp). He was the Coordinator of the Hyperspectral Imaging Network, a European project designed to build an interdisciplinary research community focused on hyperspectral imaging activities. He is the author or coauthor of more than 350 publications on remotely sensed hyperspectral imaging, including more than 90 journal citation report papers (45 since January 2011), around 20 book chapters, and over 230 conference proceeding papers. His research interests include remotely sensed hyperspectral imaging, pattern recognition, signal and image processing, and efficient implementation of large-scale scientific problems on parallel and distributed computer architectures. Dr. Plaza has guest edited seven special issues on scientific journals on the topic of remotely sensed hyperspectral imaging. He has been a Chair for the IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (Whispers 2011). He has also been serving as a Chair for the SPIE Conference on Satellite Data

Compression, Communications, and Processing, since 2009, and for the SPIE Europe Conference on High-Performance Computing in Remote Sensing, since 2011. He has been a recipient of the recognition of Best Reviewers of the IEEE Geoscience and Remote Sensing Letters in 2009 and a recipient of the recognition of Best Reviewers of the IEEE Transactions on Geoscience and Remote Sensing in 2010, a journal for which he has served as Associate Editor since 2007 and for which he has reviewed more than 260 manuscripts. He is currently serving as the Director of Education Activities for the IEEE Geoscience and Remote Sensing Society (GRSS), and as President of the Spanish Chapter of IEEE GRSS. He has been appointed Editor of the IEEE Transactions on Geoscience and Remote Sensing, with a three-year term starting on January 2013.

List of figure captions

Fig. 1. On-board processing systems scenario

Fig. 2. General view of the classical FPGA architecture

Fig. 3. Performance vs. flexibility comparison between GPPs, ASICs, and FPGAs

Fig. 4. Data Flow for Hyperspectral Image Compression

Fig. 5. Hyperspectral unmixing chain.

Fig. 6. Graphical interpretation of the N-FINDR algorithm in a 3-D space. (a) N-FINDR initialized randomly ($p = 4$). (b) Endmember replacement. (c) Final volume estimation by N-FINDR.

Fig. 7. Pseudocode of ISRA algorithm for unmixing one hyperspectral pixel vector x using a set of E endmembers

Fig. 8. Hardware architecture to implement the endmember extraction step.

Fig. 9. Hardware architecture to implement the N-FINDR algorithm

Fig. 10. Hardware architecture of the $\text{abs}(\det)$ module.

Fig. 11. Hardware architecture to implement the fractional abundance estimation step

Fig. 12. Hardware architecture to implement the ISRA module.

Figure 1

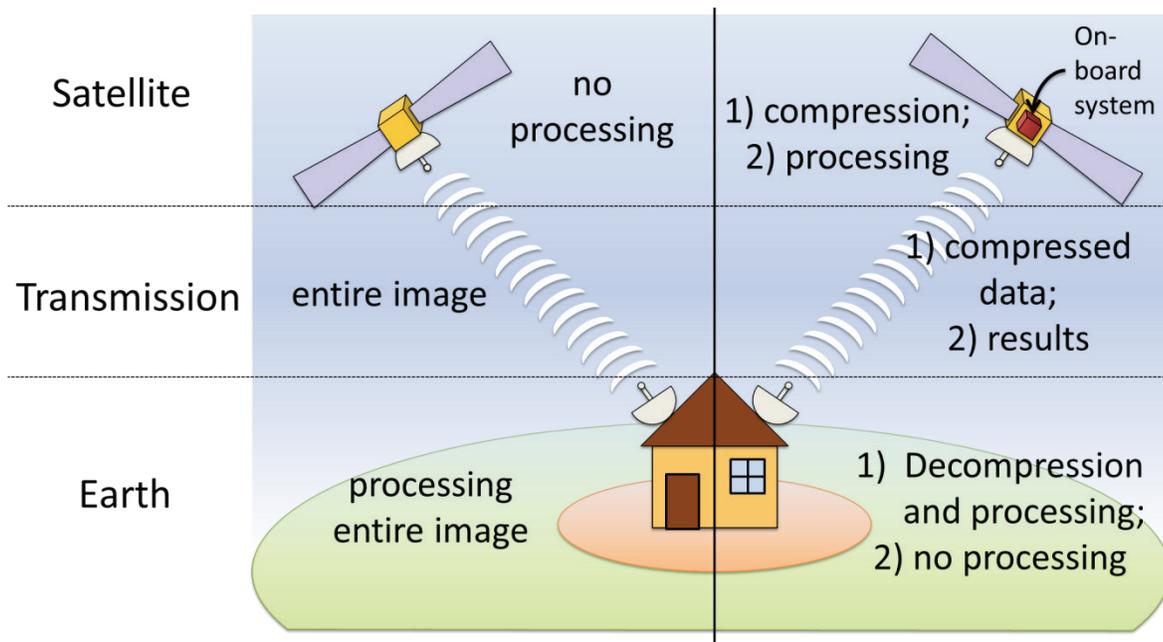


Figure 2

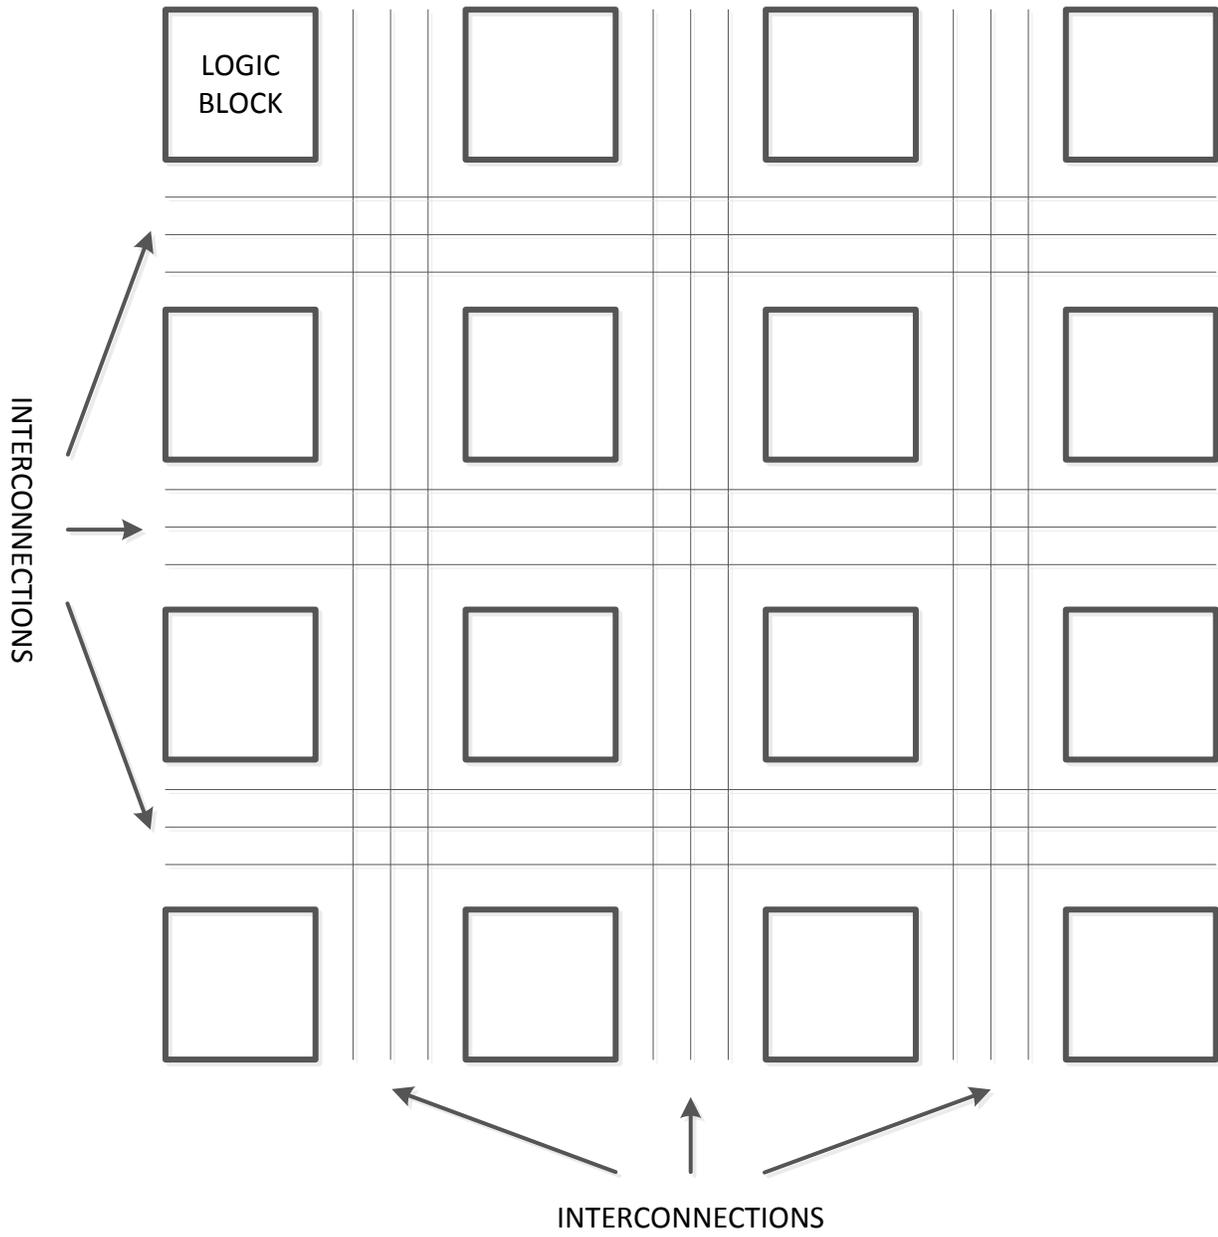


Figure 3

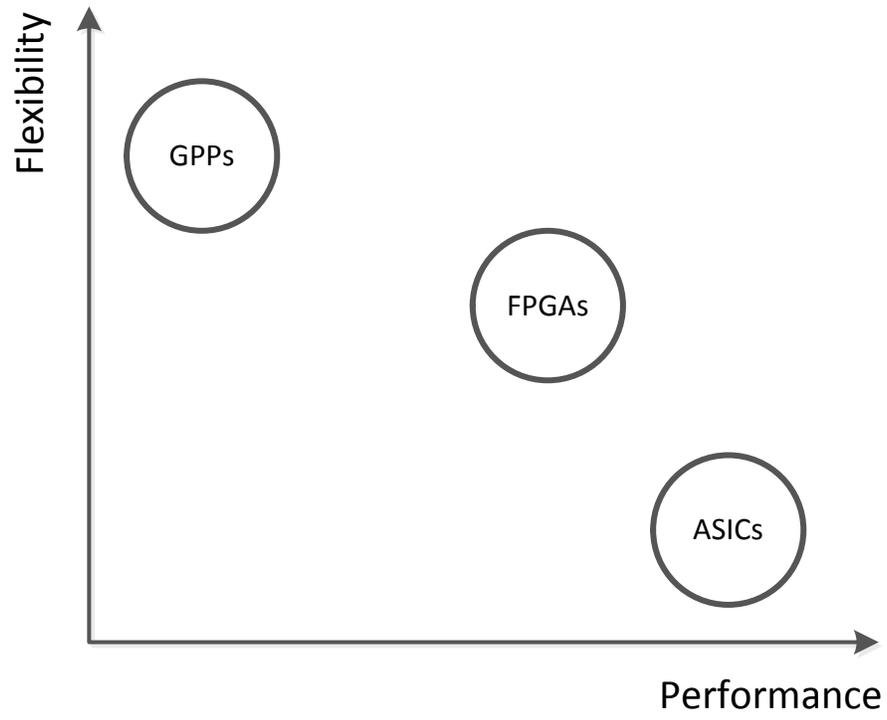


Figure 4

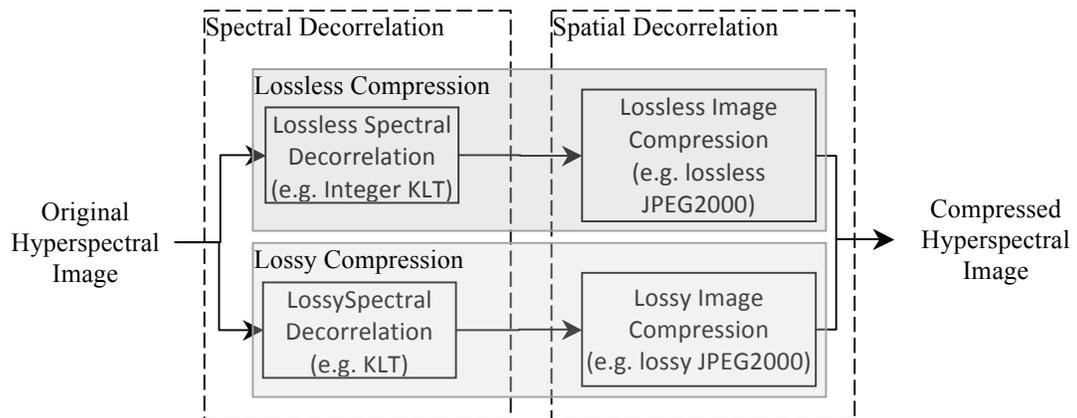


Figure 5

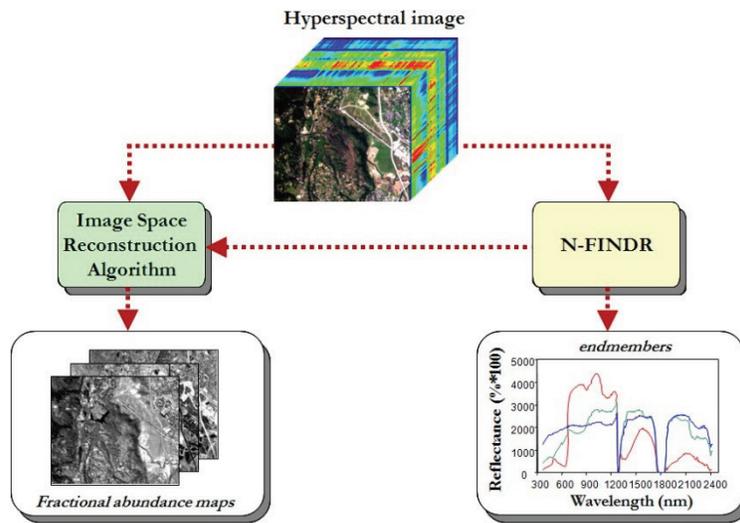


Figure 6

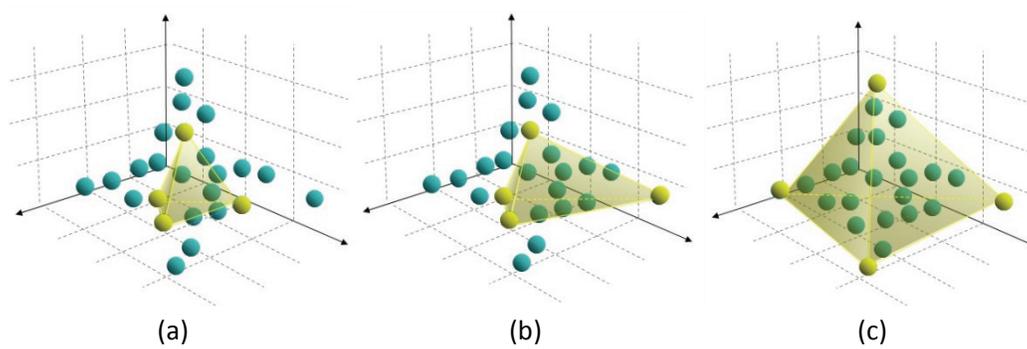


Figure 7

```

// For a certain number of iterations
for (k = 0; k < iters; k++) {
  // For all endmembers
  for (j = 0; j < p; j++) {
    // For all bands
    for (i = 0; i < n; i++) {
      // Calculate the numerator
      numerator = numerator + E[i][j] * x[i];
      // Calculate the denominator
      for (s = 0; s < p; s++) {
        dot = dot + E[i][s] *  $\Phi$ [s];
      }
      denominator = denominator + dot * E[i][j];
      dot = 0;
    }
    // Calculate the new abundance vector  $\Phi$ 
     $\Phi$ [j] =  $\Phi$ [j] * (numerator / denominator);
    Numerator = 0;
    denominator = 0;
  }
}

```

Figure 8

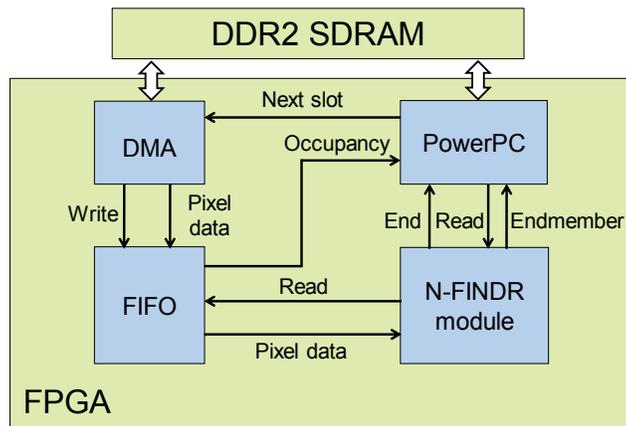


Figure 9

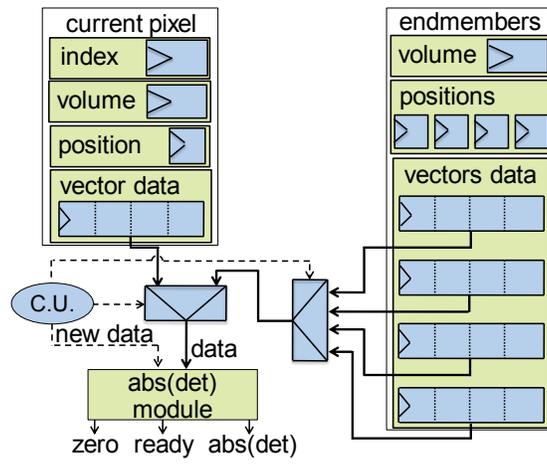


Figure 10

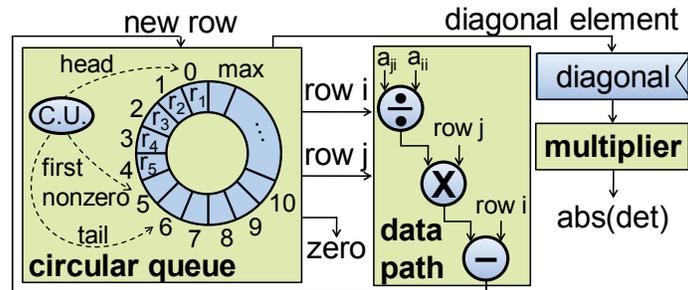


Figure 11

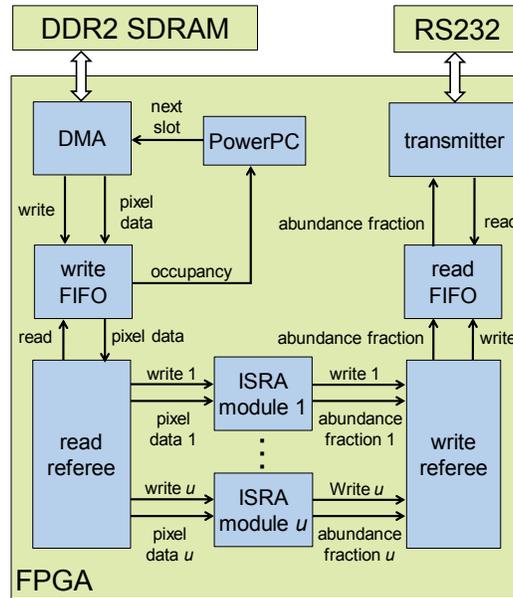


Figure 12

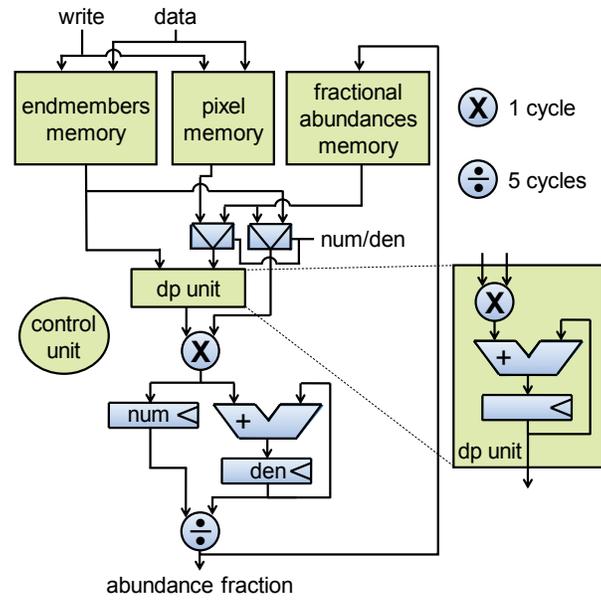


TABLE I
KLT Computational Requirements

Operations Process	Additions	Subtractions	Multiplications	Divisions	Trigonometric operations
BandMean	$M \times L \times n$	-	-	n	-
MeanSub	-	$M \times L \times n$	-	-	-
Covariance	$0.5 \times M \times L \times n \times (n+1) + M \times L \times n$	n^2	$0.5 \times M \times L \times n \times (n+1) + 0.5n(n+1) + n^2$	$n^2 + n$	-
Eigen	$5(n-1)(n-2) \times (4+2n)$	$5(n-1)(n-2) \times (4+2n)$	$5(n-1)(n-2) \times (14+8n)$	$10(n-1) \times (n-2)$	$15(n-1) \times (n-2)$
Eigen ×MeanSub	$M \times L \times n(n-1)$	-	$M \times L \times n^2$	-	-
Overall	$1.5 \times M \times L \times n \times (n+1) + 5(n-1)(n-2) \times (4+2n)$	$M \times L \times n + n^2 + 5(n-1)(n-2) \times (4+2n)$	$M \times L \times n(1.5n + 0.5) + (n+1)(0.5n+1) + n^2 + 5(n-1)(n-2) \times (14+8n)$	$11n(n-1) + 20$	$15(n-1) \times (n-2)$

TABLE II

Execution Times (in seconds) of the KLT Design for a Cluster of 32 Bands (AVIRIS Cuprite Scene)

	<i>PC</i>	<i>Cortex</i>	<i>App 1</i>	<i>Improvement %</i>	<i>App2</i>	<i>Improvement %</i>
BandMean	0.22	2.94	2.94	-	2.94	-
MeanSub	0.90	4.12	4.12	-	4.12	-
Covariance	13.80	211.65	144.74	31.6%	113.60	46.3%
Eigenvector	3.50	3.62	1.71	52.9%	3.62	-
Eigen × SubMean	16.80	445.82	292.74	34.3%	181.22	59.4%
Overall KLT	35.20	668.15	446.25	33.2%	305.40	54.3%

TABLE III
Power Consumption of the KLT Design (AVIRIS Cuprite Scene)

	<i>Static</i>	<i>Dynamic</i>	<i>Total</i>
Approach 1	8.99 mW	204.59 mW	213.58 mW
Approach 2	8.99 mW	215.4 mW	224.39 mW

TABLE IV
Hardware Resources Required by the KLT Implementations (AVIRIS Cuprite Scene)Design

		<i>Used</i>	<i>Total</i>	<i>Percentage</i>
Approach 1	FPGA Fabric (System Gates)	4248	4608	92%
	Embedded SRAM (Blocks)	4	8	50%
Approach 2	FPGA Fabric (System Gates)	3780	4608	82%
	Embedded SRAM (Blocks)	4	8	50%

TABLE V
Execution Time (in seconds) of the Integer KLT Design (AVIRIS Cuprite Scene)

<i>Operations</i>	<i>Cortex CPU</i>	<i>SoC</i>	<i>Improvement %</i>
BandMean	2.943	2.943	-
MeanSub	4.117	4.117	-
Covariance	211.65	113.6	46.3%
Eigenvector	3.625	3.625	-
PLUS	3.93	3.93	-
Lifting	562	309 *	55%
Overall Integer KLT	787.6	437.2	44.5%

*preliminary results

TABLE VI

N-FINDR and ISRA Computational Requirements (Single-precision floating point operations)

Operations Process	Additions	Subtractions	Multiplications	Divisions
N-FINDR	-	$M * L * p * \sum_{j=1}^{p-1} j$	$M * L * p * \left(\sum_{j=1}^{p-1} j + p \right)$	$M * L * p * \sum_{j=1}^{p-1} j$
ISRA	$M * L * i * (n + p * n)$	-	$M * L * i * (n + (p + 1) * n + 1)$	$M * L * i$

TABLE VII

Summary of resource utilization for the FPGA-based implementation of the N-FINDR algorithm on a VIRTEX-4 XC4VFX60 FPGA

Component	Number of endmembers (p)	Number of DSP48Es	Number of RAMB16s	Number of slice flip flops	Number of 4 input LUTs	Number of slices	Percentage of total	Maximum operation frequency (MHz)
N-FINDR module	21	128	217	20077	34155	24622	97.40	51.3
DMA Controller	-	0	0	170	531	367	1.45	102

TABLE VIII

Summary of resource utilization for the FPGA-based implementation of the ISRA algorithm on a VIRTEX-4 XC4VFX60 FPGA

Component	Number of modules (u)	Number of DSP48Es	Number of RAMB16s	Number of slice flip flops	Number of 4 input LUTs	Number of slices	Percentage of total	Maximum operation frequency (MHz)
Parallel ISRA module	16	128	54	5532	42047	22773	90.08	53.4
RS232 Transmitter	-	0	0	69	128	71	0.28	208
DMA Controller	-	0	0	170	531	367	1.45	102

TABLE IX
Power Consumption of the N-FINDR and ISRA Designs

	<i>Static</i>	<i>Dynamic</i>	<i>Total</i>
N-FINDR	455 mW	350 mW	805 mW
ISRA	455 mW	362 mW	817 mW

TABLE X

Execution times measured for the FPGA hardware implementation (HW) and for an equivalent software (SW) version

	AVIRIS Cuprite subscene (16 endmembers and 10 iterations)		EO-1 Hyperion Cuprite scene (21 endmembers and 10 iterations)	
	HW	SW	HW	SW
Endmember extraction (N-FINDR)	13.46 s	476.34 s	239.11 s	8464,99 s
Dynamic reconfiguration	694 ms	0 ms	694 ms	0 ms
Fractional abundance estimation (ISRA)	1.32 min	5.81 min	30.21 min	173.4 min
Total	1.56 min	13.75 min	34.21 min	314.48 min