

Advancing NASA's On-Board Processing Capabilities with Reconfigurable FPGA Technologies

Paula J. Pingree
*Jet Propulsion Laboratory, California Institute of Technology
United States of America*

1. Introduction

Future NASA missions will require measurements from high data rate active and passive instruments. Recent internal studies at NASA's Jet Propulsion Laboratory (JPL) estimate approximately 1-5 Terabytes per day of raw data (uncompressed) are expected, for example, from spectroscopy instruments. Implementations of on-board processing algorithms to perform lossless data reduction are required to drastically reduce data volumes to within the downlink capabilities of the spacecraft and existing ground stations. Reconfigurable Field Programmable Gate Arrays (FPGAs) such as the Xilinx™ Virtex-4 and Virtex-5 series devices can include dual core PowerPC processors thereby providing a flexible hardware and software co-design architecture to meet the on-board processing challenges of these missions while reducing the essential resources of mass and volume of earlier generation flight-qualified computing platforms such as the BAE Rad750 single board computer (SBC). Reconfigurable FPGAs also offer unique advantages over one-time programmable (OTP) FPGAs with flexible prototype development platforms that provide an important "path-to-flight" for spaceborne instruments. Reconfigurable FPGA technologies also provide in-flight flexibility with the ability to update processing algorithms as needed post-launch.

This chapter will discuss these comparative technologies and present the benefits of commercially available FPGA development platforms from Xilinx for the development of NASA's future on-board processing capabilities. Additionally, commercially available tools such as Impulse C™ have been used to adapt legacy C-code into Verilog or VHDL for implementation in FPGA fabric to achieve hardware acceleration. Key features and requirements of future NASA missions proposed within the National Research Council's Decadal Survey will be described with ideas on how these reconfigurable FPGA technologies and development tools can combine to achieve breakthrough on-board processing performance to meet their science objectives. To provide specific demonstrations of these ideas, three unique and recent design implementations on the Xilinx V4FX60 and Virtex-5 FPGAs targeted to enable future NASA missions will be presented. They include on-board processing algorithms for a) Support Vector Machine (SVM) Classifiers similar to those in operation on the Earth Observing 1 (EO-1) Hyperion instrument, b) a Fourier transform infrared (FTIR) spectrometer, and c) a new Multiangle Spectropolarimetric

Source: Aerospace Technologies Advancements, Book edited by: Dr. Thawar T. Arif,
ISBN 978-953-7619-96-1, pp. 492, January 2010, INTECH, Croatia, downloaded from SCIYO.COM

Imager (MSPI). Finally, the important issue of radiation effects from single event upsets (SEUs) and necessary mitigation schemes for SRAM-based reconfigurable FPGA devices will be addressed particularly as it depends on science data return requirements for both Earth and planetary missions.

2. Earth science objectives & the need for advanced on-board processing

Many of the missions proposed within the National Research Council's Decadal Survey (NRC, 2007) will require measurements from high data rate active and passive instruments. Internal studies at JPL estimate approximately 1-5 Terabytes of expected raw (uncompressed) data per day from spectroscopy instruments. When lossless compression algorithms are applied, this volume can be reduced to 100s of Gigabytes per day, but many of the Decadal Survey missions actually carry multiple instruments. Implementations of on-board processing algorithms to perform science data reductions can dramatically reduce the data rates, but there are additional advantages. An on-board processing platform of sufficient capability can simplify the instrument design itself leading to reductions in mass, power, and volume for missions where multiple instrument components necessitate optimization of these resources. In this way, strategic application of information technology advances can lead to measurable improvements in both instrument data reduction and design. This chapter will highlight two Decadal Survey Missions for which reconfigurable FPGA technologies are targeted to meet the on-board processing requirements of the key instrument payloads. Both the Geostationary Coastal and Air Pollution Events (GEO-CAPE) and Aerosol-Cloud-Ecosystem (ACE) missions are recommended for launch in the 2013-2016 timeframe.

The GEO-CAPE mission will analyze atmospheric gas columns for air quality forecasts and analyze ocean color for coastal ecosystem health and climate emissions (Sander et al., 2008). The key instruments for GEO-CAPE are a high-spatial resolution hyperspectral spectrometer, a low-spatial resolution imaging spectrometer and an IR correlation radiometer. The Panchromatic Fourier Transform Spectrometer (PanFITS), a new project within the NASA Instrument Incubator Program, will measure all of the trace species called out in the Decadal Survey for GEO-CAPE. With continuous sensitivity from 0.26 to 15 micron and high spectral resolution, PanFITS combines the functionality of separate UV, visible and IR instruments in a single package. These capabilities also permit PanFITS to meet the requirements for high spatial resolution hyperspectral imaging of the coastal zone. Two key enabling technologies under development for PanFITS are high-speed, high-dynamic range CMOS hybrid focal plane arrays (FPAs), and parallel, co-aligned optical trains for the ultraviolet-Visible-Near-infrared (UV-Vis-NIR) and mid-IR bands. Section 4 of this chapter describes the role of reconfigurable FPGAs in the development of the PanFITS instrument to meet the on-board processing requirements for the GEO-CAPE mission.

The ACE mission will analyze aerosol and cloud profiles for climate and water cycles and analyze ocean color for open ocean biogeochemistry. The Decadal Survey identifies a multi-angle, multispectral, high-accuracy polarization imager as a key requirement for the ACE mission. JPL has been developing a Multiangle SpectroPolarimetric Imager (MSPI) as a candidate to fill this need. Using passive remote sensing to determine aerosol and cloud microphysical parameters necessitates a variety of constraints to overcome retrieval non-uniqueness, and places stringent design requirements on a cohesive, next-generation passive satellite imager. Fusion of ultraviolet (UV) to SWIR multispectral, multiangular, and polarimetric observations is important because each brings sensitivity to different aspects of

particle spatial distributions and microphysics (Yu et al., 2006; Kahn et al., 2004). UV wavelengths are sensitive to aerosol absorption and height, visible/near-infrared (VNIR) wavelengths to fine mode aerosol size distributions, and SWIR wavelengths to coarse mode aerosol and cloud particle sizes. Multiangular radiances distinguish particle shapes, and improve sensitivity to optical depth, notably over bright surfaces. Polarimetry is the only means of determining particle real refractive index, providing compositional information. Moderately high spatial resolution resolves intercloud scales and aerosol gradients in urban settings. A broad swath is important for providing environmental context to aerosol and cloud spatial relationships, distinguishing airmass types, validating transport models, and overflying field campaigns and suborbital instruments. The technological design considerations of integrating all of these attributes into a single instrument are compounded by the need to acquire accurate multispectral intensity imagery (3% uncertainty) simultaneously with accurate degree of linear polarization (DOLP) imagery (0.5% uncertainty). The MSPI technology development plan has been systematically reducing risk in critical areas, with consequent direct benefit to the ACE mission. An essential part of this plan is on-board processing for digital signal handling. The role of reconfigurable FPGAs for MSPI on-board processing is discussed further in Section 4 of this chapter.

3. Comparative technologies

On-board computation has become a bottleneck for advanced science instrument and engineering capabilities. Currently available spacecraft processors have high power consumption, are expensive, require additional interface boards, and are limited in their computational capabilities. Recently developed hybrid FPGAs, such as the Xilinx Virtex-4FX60 and Virtex-5, offer the versatility of running diverse software applications on embedded processors while at the same time taking advantage of reconfigurable hardware resources, all on the same chip package. These tightly coupled hardware/software co-designed systems are lower power and lower cost than general-purpose SBCs, and promise breakthrough performance over radiation-hardened SBCs, leading to a new architecture for future spaceborne instrument development.

Computational Platform	Performance (DMIPS)
RAD750 SBC	240
Xilinx Virtex-II Pro	450
Xilinx Virtex-4	680
Xilinx Virtex-5	> 1000

Table 1. Performance: SBC vs. Embedded FPGAs (<http://www.xilinx.com>)

The Rad750 SBC has flight-proven heritage serving as the primary spacecraft computer for JPL's 2005 Deep Impact mission (<http://deepimpact.jpl.nasa.gov>). For future science instruments, however, baselining a Rad750 SBC is likely to exceed the instrument's allocations for mass and volume, especially when there are multiple science instruments to be accommodated on a single spacecraft platform. In some cases a single Rad750 SBC may not even meet the science instruments requirements for on-board data processing as was the case for an instrument called MATMOS that was proposed for the 2007 Mars Scout mission

opportunity (See Section 4.2). MATMOS is a precursor instrument that led to the development of PanFTS targeted for the GEO-Cape mission. The on-board processing FPGA development started on MATMOS has been extended toward a baseline architecture for PanFTS using the Xilinx Virtex-5 FPGA with embedded PowerPC440 processors.

The advantage reconfigurable FPGAs offer in flexibility over OTP FPGAs in the instrument development and operations cycles can be illustrated by the following timeline. NASA's upcoming Juno mission to Jupiter is scheduled to launch in August 2011 and arrive at Jupiter in 2016 for a nominal one year mission orbiting the planet. A key instrument on the Juno spacecraft is the Microwave Radiometer (MWR), under development at JPL (Pingree et al., 2008-a). The MWR Electronics Unit contains a single OTP FPGA for the command and telemetry interface to the spacecraft as well as to control the instrument's six (6) receivers. To meet the MWR instrument development schedule, the final FPGA design was programmed and installed on the flight electronics board in June 2009. The implication is that 2+ years before the spacecraft is launched the MWR FPGA design cannot be modified, improved or corrected without re-working the flight board at a significant impact to the project's cost and schedule. Furthermore, there is an additional 5 years of cruise to Jupiter that will include MWR instrument calibration activities, and still the FPGA design remains fixed. Irregardless of potential in-flight findings where updating the FPGA may be desirable to improve the instrument's capabilities to better meet the science objectives of the mission, doing so is impossible with a OTP FPGA. This would not be the case with an SRAM-based reprogrammable FPGA such as the Xilinx Virtex-5. Especially with the future trend for increasing complexity in on-board processing requirements, it is particularly beneficial to have the flexibility to improve and optimize algorithms throughout the instrument development and operations cycles.

Other JPL missions have used Xilinx reprogrammable FPGAs. Four Xilinx XQR4062XL FPGA devices controlled the Mars Exploration Rovers (MER) lander pyrotechnics that were crucial to the Entry, Descent and Landing phase of the mission. The MER Spirit and Opportunity vehicles, still in operation after over 5 years on the surface of Mars, contain the Xilinx XQVR1000 FPGA in their motor controller boards. The Mars Reconnaissance Orbiter (MRO) Electra software radio uses the Xilinx VirtexII-8000 FPGA. Each of these devices, however are not hybrid FPGAs; they do not contain embedded PowerPC processors. The Xilinx hybrid FPGAs such as the V4FX60 and Virtex-5 showcased in the next Section's design cases do not have space-flight heritage and therefore represent the cutting edge of computing platforms required to meet the requirements of the Decadal Survey missions.

4. Recent on-board processing FPGA design cases

To provide specific demonstrations of the ideas and advantages of reprogrammable FPGAs that have been discussed, three unique and recent design implementations on the Xilinx V4FX60 and Virtex-5 FPGAs targeted to enable future NASA missions are presented. They include on-board processing algorithms for a) Support Vector Machine (SVM) Classifiers similar to those in operation on the EO-1 Hyperion instrument, b) a Fourier transform infrared (FTIR) spectrometer, and c) a new Multiangle Spectropolarimetric Imager (MSPI).

4.1 SVMs for hyperspectral classification

Fast and accurate on-board classification of image data is a critical part of modern satellite image processing. For Earth sciences and other applications, space-based payloads make use

of intelligent, machine-learning algorithms and instrument autonomy to detect and identify natural phenomena such as flooding, volcanic eruptions, and sea ice break-up. JPL has developed support vector machine (SVM) classification algorithms used on board spacecraft to identify high-priority image data for downlinking to Earth. These algorithms also provide onboard data analysis to enable rapid reaction to dynamic events. These onboard classifiers help reduce the amount of data downloaded to Earth, greatly increasing the science return of the instrument.

SVM classification algorithms are flying today using computational platforms such as the RAD6000 and Mongoose V processors; these SBCs are of an even older generation than the Rad750 discussed in Section 3. These legacy processors have only limited computing power, extremely limited active storage capabilities, and are no longer considered state-of-the-art. For this reason, onboard classification has been limited to only the simplest functions running on only a subset of the full instrument data: for example, only 11 of 242 image bands are classified in the case of the Hyperion instrument on the EO-1 satellite. FPGA coprocessors are an ideal candidate for future implementations of these algorithms to provide significant improvement in onboard classification capability and accuracy when compared to the legacy processing platforms now flying.

SVMs are also well-suited to onboard autonomy applications. The property that makes SVMs particularly applicable is the asymmetry of computational effort in the training and testing stages of the algorithm. Classifying new data points requires orders of magnitude less computation than training because the process of training a SVM requires solving a quadratic optimization problem. SVM training requires $O(n^3)$ operations, where n is the number of training examples. In contrast, testing a new vector with a trained SVM requires only $O(n)$ operations. Faster algorithms that exploit the specific structure of the SVM optimization problem exist (Platt, 1998), but the training remains the primary computational bottleneck.

After a SVM is trained, many of the weights, w_i , will be equal to zero. This means that these terms can be ignored in the classification formula. Those input vectors that have a corresponding non-zero weight are called *support vectors*. Even more computational savings can be realized in the case of using a linear kernel function. The weighted sum over the kernel function is associative, so all the support vectors can be collapsed into a single vector with a single weight. Reducing the number of support vectors is key to successfully deploying a SVM classifier onboard a spacecraft where there are severe constraints on the available computational resources. Previously deployed classifiers (Castano et al., 2006) have used such reduced-set methods, but were still constrained to operate on only a subset of the available classification features. Removing such bottlenecks is critical to realizing the full potential of SVMs as an onboard autonomy tool.

In 2007, a small JPL team evaluated the effectiveness of FPGAs for SVM algorithms for a legacy snow-water-ice-land (SWIL) classifier, originally developed for the Hyperion instrument, on the Xilinx ML410 development board that includes the Virtex-4 FX60 FPGA (Pingree et al., 2008-b). For rapid development, the Impulse CTM autocode toolset, provided by Impulse Accelerated Technologies, was used to adapt legacy C-code into HDL¹ for implementation in the FPGA fabric. The Impulse C compiler generates synthesizable HDL files ready to use with the Xilinx EDK tools. In addition to generating HDL files, the Impulse

¹ High-level Design Language (HDL)

compiler also generates additional files required by the EDK tools, including PLB (Processor Local Bus) and APU (Auxiliary Processor Unit) bus interfaces. The Impulse C compiler performs a variety of low-level optimizations, including C statement scheduling and loop pipelining, saving application developers a great deal of time that would otherwise be spent performing tedious, low-level hardware optimization. The Impulse C compiler performs these optimizations and generates hardware in the form of either VHDL or Verilog. This hardware can then be synthesized using FPGA tools such as Xilinx ISE™ software and Platform Studio.

On the processor side, the compiler generates run-time libraries ready for use on the embedded PowerPC processor. While it is arguable that auto-code generation tools may not produce Verilog or VHDL that is as efficient as a programmer may develop directly by hand, but for already-existing legacy C-code algorithms that are intended for acceleration by FPGA implementation, the Impulse C tools can provide an efficient path toward rapid prototype development. This approach to implementing the Hyperion linear SVM on the Virtex-4 FX60 FPGA, as well as additional experiments that were performed using an increased number of data bands and a more sophisticated SVM kernel, provided rapid results to show the potential for more efficient, higher performance onboard classification using FPGA-embedded algorithms.

When using FPGAs with embedded processors, efficient partitioning of algorithms between software and hardware is important to achieve high performance. The image file input and classification file output are managed within the FPGA's embedded PowerPC processor. Figure 1 shows an example architecture for this FPGA-based implementation.

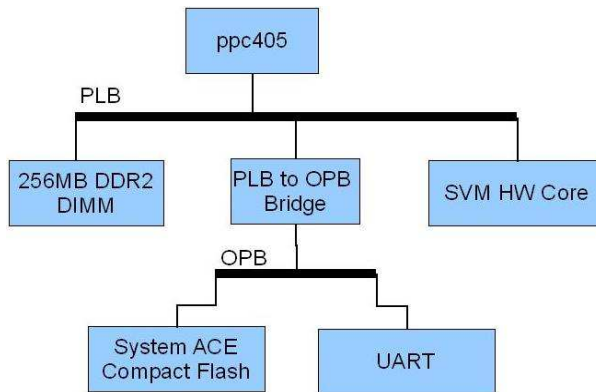


Fig. 1. SVM Classifier FPGA Architecture (Pingree et al., 2008-b).

In support of partitioned software/hardware applications such as this, the Impulse tools include a library of C-compatible functions that implement a number of process-to-process communication methods. These methods include streaming, shared memory, and message passing. For this application, the Impulse C streaming programming model was the obvious choice. In Impulse C streaming applications, hardware and software processes communicate primarily through buffered data streams implemented directly in hardware. This buffering of data makes it possible to write parallel applications at a relatively high level of abstraction without the cycle-by-cycle synchronization that would otherwise be required. Impulse C functions are used to open and close data streams, read or write data on the streams, and, if desired, send status messages or poll for results. In the case of the

Virtex-4FX, stream reads and writes can be specified as operations that take advantage of either the PLB or the auxiliary peripheral unit (APU) interface.

The software side of the SVM application is coded in C and compiled to the embedded PowerPC 405 processor using the Xilinx EDK tools. The embedded software application reads an input image file consisting of 857,856 pixels (based on actual Hyperion image size). The software-side application streams the image data to the SVM, which is also written in C but has been compiled (using the Impulse C-to-FPGA compiler) to FPGA hardware. The SVM hardware process performs the required SVM operation on the image and streams the results back to the PowerPC 405 processor. The processor then writes the pixel classifications (e.g., snow, water, ice, land, cloud, or unclassified) to an output file.

The Impulse C-converted HDL code is synthesized for the Virtex-4FX FPGA using the Xilinx ISE & EDK development environment. Synthesis of this SVM application results in the following resource utilization for the V4FX60 device on the ML410 development platform (Table 2). As shown, the FPGA can easily accommodate the algorithm that was severely constrained within the EO-1 Hyperion software implementation on the Mongoose processor.

FPGA RESOURCES	V4FX60 (on ML410)
Number of Slices:	1151 out of 25280 (4%)
Number of Slice Flip Flops:	1290 out of 50,560 (2%)
Number of 4 input LUTs:	1838 out of 50560 (3%)
Number of FIFO/RAMB16s	2 out of 232 (1%)
Number of DSP48s:	4 out of 128 (3%)

Table 2. Impulse C Resource Report for SVM Application.

Post-processed image visualizations of the Virtex-4FX60 FPGA implementation of the SVM classification algorithm show excellent agreement with the results from EO-1 (Figures 2 & 3). Having successfully implemented the legacy SVM designed for Hyperion with ample FPGA resources available, extensions to the C-language algorithm were applied to increase overall system performance for on-board classification. Recall that in the Hyperion instrument algorithm currently in use on the EO-1 spacecraft, a linear SVM kernel operates on 11 of 242 available bands of pixel data. Future implementations for SVM classification with FPGAs can consider operating on a larger number of bands for a more complete classification set and/or consider using a nonlinear kernel for improved classification accuracy. Using the C-to-hardware compiler enables quick experimentation with alternative implementations such as those shown in Figure 4.

The Hyperion instrument, as part of the NASA New Millennium Program Earth Observing - 1 Mission, was the first imaging spectrometer to routinely acquire science grade data from Earth orbit (Pearlman, 2003). EO-1 was launched on a Delta 7320 rocket from Vandenberg Air Force Base on November 21, 2000. The Hyperion instrument was built by TRW, Inc. (now Northrop Grumman Space Technology) and was designed to provide high quality calibrated data for hyperspectral application evaluations (Pearlman et al., 2001). Hyperion had a single telescope and two spectrometers: one visible/near infrared (VNIR) spectrometer and one short-wave infrared (SWIR) spectrometer producing 242 bands in the range of 400-2500 nm with 10 nm bandwidths. With Hyperion, each pushbroom image frame captured the spectra from an area 30 m along-track by 7.7 km cross-track that was




Fig. 2. A comparison of the results from A) the FPGA SVM implementation, B) the legacy C-code SVM, and C) the original hyperspectral image. The color key is blue = water, cyan = ice, dark purple = snow, lavender = unclassified.




Fig. 3. The black pixels in image (B) indicate the indices where the classifications of pixels (A) and (C) were not identical.




Fig. 4. A comparison of the results from A) the 11-band linear SVM hardware implementation, B) a 30-band linear SVM hardware implementation, C) a non-linear (polynomial) SVM hardware implementation, to D) the original EO-1 hyperspectral image. Color differences between image B) and the other 2 images (A & C) are due to the different bands (qty. 11 vs. 30) selected for each classification.

sufficient to address most land cover issues. The forward motion of the satellite created a sequence of frames that were combined into a two-dimensional spatial image with a third dimension of spectral information (called a "3-d data cube").

In January 2004, as part of the New Millennium Program ST-6 Autonomous Sciencecraft Experiment (ASE), autonomy software has been executing on the EO-1 Mongoose V processor using Hyperion image data for on-board science analysis and autonomous planning and execution of instrument observations (<http://ase.jpl.nasa.gov>). ASE employs four unique classifiers on Hyperion image pixel data. One is the SWIL classifier described previously in this Section. Following pixel classification, ASE autonomy software triggers new observations based on the analysis of the current observation. This is known as event detection. Specific events of interest on EO-1 include sea ice break-up and formation, lake freeze and thaw, and active volcanism. The limitations of executing the software SWIL classifier on the Mongoose V processor are clear. Of the 242 available bands of Hyperion image data only 11 bands can be processed.

Accommodating reprogrammable FPGAs in the architecture and design of future NASA instruments can provide a highly capable and flexible computing platform for on-board science data processing. The Decadal Survey's Hyperspectral Infrared Imager (HypIRI) Mission is poised to benefit from reprogrammable FPGAs. The HypIRI mission has heritage in the EO-1 Hyperion imaging spectrometer as well as in JPL's Moon Mineralogy Mapper (M³) instrument on the Indian Moon-orbiting mission, Chandrayaan-1. Using imaging spectroscopy, the HypIRI mission would obtain global coverage from LEO with a repeat frequency of 30 days at 45-m spatial resolution. A pointing capability is required for frequent and high-resolution imaging of critical events, such as volcanos, wildfires and droughts. While there is no present plan to include on-board classifiers such as the SWIL classifier on HypIRI, a cloud detection algorithm implemented in a reconfigurable FPGA is under evaluation for the mission (LeMoigne, 2007).

4.2 FTIR Spectrometer

The Panchromatic Fourier Transform Spectrometer (PanFTS) is a NASA Instrument Incubator Program (IIP) funded development to build and demonstrate a single instrument capable of meeting or exceeding the requirements of the GEO-CAPE mission. An FTS is a Michelson interferometer in which the optical path difference of light rays is continuously varied with moving mirrors. Using photovoltaic detectors, this modulated light is converted to an electric signal known as an interferogram. For broadband sources, the interferogram exhibits a large peak at zero path difference where all wavelengths interfere constructively. The PanFTS development effort has strong heritage in the MATMOS FTIR² spectrometer instrument concept designed to measure the Mars atmospheric composition using solar occultation from orbit as proposed (but not selected) for the 2007 Mars Scout mission. The MATMOS instrument would measure the 850-4300 *cm*⁻¹ region of the IR spectrum of sunlight as it shines through the Martian atmosphere. MATMOS would record roughly 26 spectra per occultation, with each containing 172,500 spectral elements. There are two occultations per orbit, one for sunrise and one for sunset. The duration of an occultation is

² Fourier Transform Infra Red (FTIR)

between 78 and 169 seconds, thus requiring that each spectrum be collected in 3.0 to 6.5 seconds (Pingree et al., 2007). The MATMOS FTS utilizes three separate detectors in the process of collecting occultation spectra. An *HgCdTe* detector is used to collect longer wavelengths (5 μm to 2 μm) and an *InSb* detector collects shorter wavelengths (5 μm to 2 μm). A *Ge* detector is used to collect the reference laser interferogram (used to measure the path difference - internal to the FTS). For each orbit, the three detectors produce 659 Mbytes of raw data that must be processed and compressed prior to downlink.

The data processing consists of five steps: interferogram resampling, phase correction, FFT (Fast Fourier Transform), spectra averaging, and lossless compression. Re-sampling converts the time-domain signal to the path difference domain, removing frequency modulation in the process and reducing the number of points for each solar detector. Phase correction makes the interferogram symmetrical about the zero path difference (ZPD), a point where the two moving mirrors inside the interferometer are at equal distance to the beam splitter. This allows the two symmetrical halves of the interferogram to be averaged together. The spectrum is then computed with an FFT, reducing the dynamic range of the interferogram thus allowing it to be represented with fewer data bits. Averaging scans taken above the atmosphere and then performing lossless compression further reduces the volume of data to be transmitted to Earth.

In a 2007 technology demonstration, the Xilinx V4FX60 FPGA was evaluated for its FTIR spectrometer data processing capability targeting the MATMOS instrument development (Bekker et al., 2008). By optimizing floating-point calculations necessary for processing and compression of MATMOS data prior to downlink, a more than 8x reduction in execution time was achieved on the FPGA however, these results still lagged behind the Rad750's processing capabilities. In 2008, the FTIR spectrometry algorithm was targeted to the most recently available Virtex-5FXT FPGA (Bekker et al., 2009). The V5FXT FPGA contains the more powerful PPC440 processor, more cache, and improved memory interfaces over those of the V4FX, as well as an improved auxiliary processor unit (APU) controller and floating-point unit (FPU). Preliminary results for the MATMOS FTIR on-board processing algorithm on the V5FXT show a nearly 5-x improvement over the V4FX implementation and execution times that now surpass the Rad750. The FTIR V5FXT system is shown in Figure 5.

The PanFITS design for the GeoCAPE mission combines atmospheric measurement capabilities in the IR and UV-VIS³ with the ability to measure ocean color by using imaging FTS to provide full spatial coverage. For the atmospheric composition, the instrument includes up to four Focal Plane Arrays (FPA) of 128 x 128 pixels that are read at a frame rate of 16 kHz. JPL has developed an interface that records pixel data from commercially available IR FPAs that are capable of the required frame rate at a lower spatial coverage. This interface uses high speed ADCs⁴ and the Xilinx Virtex-5FXT FPGA. FTIR on-board processing development efforts on the V5FXT FPGA continue at JPL for the PanFITS instrument.

³ Ultra Violet Visible (UV-VIS)

⁴ Analog to Digital Converter (ADC)

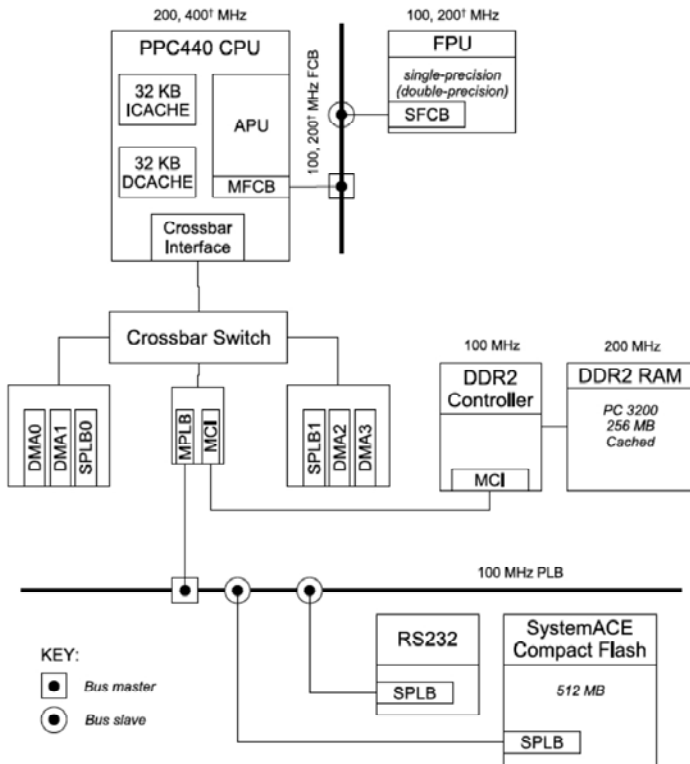


Fig. 5. V5FXT FTIR system with FPU co-processor, shown in multiple configurations (Bekker et al., 2009). † indicates an alternate configuration.

4.3 The Multi-angle Spectro-Polarimetric Imager (MSPI)

The Multi-angle Spectro-Polarimetric Imager (MSPI) is an advanced instrument concept in development at JPL to produce a highly accurate multi-angle, multi-wavelength polarimeter to measure cloud and aerosol properties as called for by the Aerosol-Cloud-Ecosystem (ACE) mission concept in the Earth Sciences Decadal Survey. The MSPI instrument will use a set of 9 cameras (8-fixed and 1-gimballed)⁵, each associated with a given along-track view angle in the 0°-70° range (see Figure 6). Each camera must eventually process a raw video signal rate around 95 Mbytes/sec over 16 channels.

The greatest challenges of the MSPI instrument are the stringent demand on degree of linear polarization (DOLP) tolerance over a wide swath, and the need to acquire polarimetric and multispectral intensity imaging simultaneously from the UV-SWIR⁶. In an attempt to achieve necessary accuracy of the DOLP of better than 0.5%, the light in the optical system is subjected to a complex modulation designed to make the overall system robust against

⁵ Number of cameras is not finalized; may be 7-9.

⁶ Ultra-Violet Short Wave Infra Red (UV-SWIR)




Fig. 6. A conceptual layout of the MSPI instrument. A set of fixed cameras view different angles and a gimballed camera provides high angular resolution for selected Earth targets and camera-to-camera calibration (Diner et al., 2007).

many instrumental artifacts that have plagued such measurements in the past. This scheme involves two photoelastic modulators that are beating in a carefully selected pattern against each other (Diner et al., 2007). In order to properly sample this modulation pattern, each of the proposed nine cameras in the system needs to read out its imager array about 1000 times per second, resulting in two orders of magnitude more data than can typically be downlinked from the satellite.

A key technology development needed for MSPI is on-board signal processing to calculate polarimetry data as imaged by each of the 9 cameras forming the instrument. With funding from NASA's Advanced Information Systems Technology (AIST) Program, JPL is solving the real-time data processing requirements to demonstrate, for the first time, how signal data at 95 Mbytes/sec over 16-channels for each of the 9 multi-angle cameras in the spaceborne instrument can be reduced on-board to 0.45 Mbytes/sec. This will produce the intensity and polarization data needed to characterize aerosol and cloud microphysical properties. The onboard processing required to compress this data involves least-squares fits of Bessel functions to data from every pixel, effectively in real-time, thus requiring an on-board computing system with advanced data processing capabilities in excess of those commonly available for space flight. A Xilinx Virtex-5 FPGA-based computing platform is currently under development at JPL to meet MSPI's on-board processing (OBP) requirements.

As a brief polarimetry imaging overview, DOLP is calculated by equation (1), where I is the total intensity, and Q and U describe linear polarization.

$$\text{DOLP} = \sqrt{(Q/I)^2 + (U/I)^2} = \sqrt{q^2 + u^2} \quad (1)$$

To achieve the high degree of accuracy in DOLP, two photo-elastic modulators (PEMs) are included in the MSPI optical path to modulate the Q and U polarization components of the Stokes vector. One full cycle of the modulated polarization signal occurs in the time of one 40-msec frame, set by the beat frequency of the two PEMs. Each cycle of the modulation must be "oversampled" to create a hi-fidelity digital representation of the polarization components. The baseline is to sample the modulation 32 times per frame – thereby creating 32 sub-frames per frame. Compared to MISR⁷ cameras, each with 4 spectral channels, the raw video data rate that must be handled by MSPI is increased by a factor of 256 (32x due to oversampling; 4x due to expansion of the number of channels, and 2x due to correlated double sampling to suppress read noise in the Si-CMOS readout). A single 16-channel MSPI camera (one of nine) must process 95 Mbytes/sec of raw video data. A computationally intensive linear least-squares algorithm must also be applied to perform data reduction for video processing of the signal output from the photo-detector array. These data reductions can be performed (without sacrificing the information content of the camera product for science) based on how the calculations for digital signal processing are implemented in the reconfigurable FPGA.

The MSPI on-board processor collects data as it streams out from the focal plane of the camera, calculates the basis function values, and computes the least-squares fit of the data using the basis functions. The result of the on-board processing is the reduction of dozens of samples acquired during a 40-msec frame to five parameters. In 2008, the Xilinx Virtex-4FX60 FPGA, including PowerPC405 processors, was used to implement a least-squares fit Bessel function fitting algorithm to generate a pixel data stream (Norton et al., 2009). The algorithm extracts intensity and polarimetric parameters in real-time thereby substantially reducing the image data volume for spacecraft downlink without loss of science information. The accuracy results of the FPGA design indicate that the OBP contribution to the MSPI degree of linear polarization (DOLP) accuracy requirement of less than 0.5% error is on the order of only 0.001%. The Virtex-4FX60 FPGA-based design for MSPI OBP is shown in Figure 7. It provides a successful prototype for the 3-channel ground-based instrument and indicates a path-to-flight for the full 16-channel space-flight instrument proposed for the ACE mission.

Current efforts to advance the MSPI OBP design target the Xilinx Virtex-5 FPGA for its advanced radiation hardness and increased performance capabilities. This 3-year task funded by the NASA ROSES AIST Program will meet the following objectives: (1) complete the design of the 16-channel polarimetric processing algorithm with migration and testing on the Xilinx Virtex-5 FPGA and development board; (2) integrate the on-board processor into the camera brassboard system; (3) perform FPGA design trades to optimize performance and explore how DSP features can be incorporated into the design; and (4) perform laboratory and airborne validation of the OBP system with real-time retrieval of polarimetry data.

⁷ Multi-angle Imaging SpectroRadiometer (MISR) is an in-flight instrument on the Terra spacecraft for the Earth Observing System (EOS) Mission, launched in 1999.

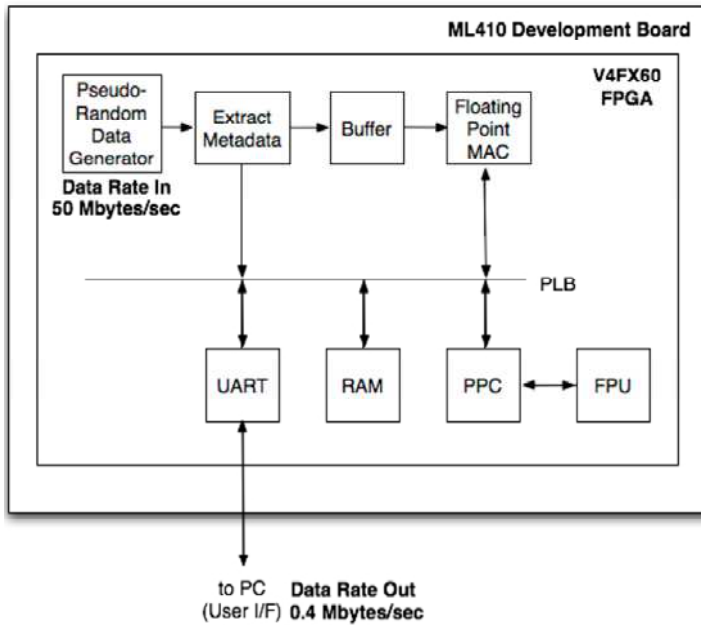


Fig. 7. Top Level Block Diagram of MSPI On-Board Processing Co-Design on the Xilinx Virtex-4 FPGA (Norton et al., 2009).

5. Radiation effects and Single Event Upset (SEU) mitigation

From NASA's Preferred Reliability Practice No. PD-ED-1258, Space Radiation Effects on Electronic Components in Low-Earth Orbit, April 1996:

"Radiation in space is generated by particles emitted from a variety of sources both within and beyond our solar system. Radiation effects from these particles can not only cause degradation, but can also cause failure of the electronic and electrical systems in space vehicles or satellites. Even high altitude commercial airliners flying polar routes have shown documented cases of avionics malfunctions due to radiation events."

"Experience with many spacecraft since Explorer I shows that higher electron concentrations are observed between 45 degrees and 85 degrees latitude in both the northern and southern hemispheres, indicating that the belts descend to a lower altitude in these regions. For low inclination orbits, less than 30 degrees, the electron concentrations are relatively low. Due to the earth's asymmetric magnetic field, a region in the Atlantic near Argentina and Brazil, known as South Atlantic Anomaly (SAA), has relatively high concentrations of electrons. The SAA is known to cause problems such as: single event upsets (SEU)."

SRAM-based reconfigurable FPGA devices are susceptible to SEUs. A necessary feature of any space-flight qualified Xilinx SRAM-based FPGA design, such as those described for

development on the Virtex-4 and Virtex-5 FPGAs, is to mitigate the effects of radiation SEUs. For Virtex-4 designs, SEU mitigation techniques such as Triple Modular Redundancy (TMR) to triplicate logic in the FPGA, presuming there are sufficient remaining resources in the device, as well as running the dual-core processors in lock-step may be employed. The Virtex-5 FPGA is advertised by Xilinx to be Rad-Hard By Design (RBDH), potentially eliminating the need for SEU mitigation techniques to be added into the design. For future low Earth-orbiting science instruments such as PanFITS (for the GEO-CAPE mission) and MSPI (for the ACE mission), the tolerance to occasional SEUs may be acceptable. The simplest approach for these instruments may be to include only SEU detection in the design and when detection occurs re-load the FPGA configuration file, a key advantage to these reconfigurable computing platforms. This is a viable strategy for non-critical applications that can withstand occasional interruption for re-configuration as may be the case for global mapping science instruments.

6. Conclusions

Hybrid or system-on-a-chip (SOC) FPGAs with embedded processors are demonstrating levels of performance and efficiency that were previously impossible using traditional processors for spaceborne computational platforms. Hardware acceleration of science instrument algorithms promises to dramatically improve onboard data processing in future NASA science missions as required by the Decadal Survey. Software-to-hardware autocode design tools can play an important role in the fast prototyping and development of legacy algorithms into hardware accelerated FPGA implementations. The Xilinx FPGA development platforms provided an excellent and cost-effective prototyping environment and a path-to-flight for future instrument on-board processing technology development. The Xilinx Virtex-4 and Virtex-5 FPGA-based developments and capabilities presented in the design cases of Section 4 respond directly to the future needs of the Decadal Survey missions for instrument science data on-board processing. The results to date demonstrate the benefits of FPGA-based processing for spectroscopy and image processing. The new RHBD architecture of the Virtex-5 FPGA promises to resolve the SRAM-based FPGA limitation of SEU susceptibility.

(Part of) This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

7. References

- Bekker, D.; M. Lukowiak, M. Shaaban, J-F. Blavier, & P. Pingree (2008). A Hybrid-FPGA System for On-Board Data Processing Targeting the MATMOS FTIR Instrument, *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2008.
- Bekker, D.; J-F. Blavier, G. Toon, & C. Servais (2009). An FPGA-Based Data Acquisition and Processing System for the MATMOS FTIR Instrument, *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2009.
- Castano, R.; N. Tang, T. Dogget, S. Chien, D. Mazzoni, R. Greely, B. Cichy, & A. Davis (2006). Onboard classifiers for science event detection on a remote sensing spacecraft, *Proceedings of the 12th ACM SIGKDD International conference of Knowledge Discovery and Data Mining*, ACM Press, (2006), 845 – 851.

- Diner, D.; A. Davis, B. Hancock, G. Gutt, R. Chipman, & B. Cairns (2007). Dual-photoelastic-modulator-based polarimetric imaging concept for aerosol remote sensing, *Applied Optics*, Vol. 46 Issue 35, pp.8428-8445 (2007).
- Kahn, R.A., J.A. Ogren, T.P. Ackerman, J. Bösenberg, R.J. Charlson, D.J. Diner, B.N. Holben, R.T. Menzies, M.A. Miller, & J.H. Seinfeld (2004). Aerosol data sources and their roles within PARAGON. *Bull. Amer. Meteorol. Soc.* 85, 1511-1522.
- LeMoigne, Jacqueline (2008). "A Reconfigurable Computing Environment for On-Board Data Reduction and Cloud Detection", 2.7 *Buckner Hypsiri Technology Investments Presentation_Final.pdf*.
- National Research Council (NRC), Committee on Earth Science and Applications from Space (2007). *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond*. The National Academies Press, Washington, DC, 437 pp.
- Norton, C.; T. Werne, P. Pingree, & S. Geier (2009). An Evaluation of the Xilinx Virtex-4 FPGA for On-Board Processing in an Advanced Imaging System, *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2009.
- Pearlman, J.; S. Carman, C. Segal, P. Jarecke, P. Barry, & W. Browne (2001). Overview of the Hyperion imaging spectrometer for the NASA EO-1 mission, *IEEE International Geoscience and Remote Sensing Symposium*, 6, pp 3504-3506, 2001.
- Pingree, P.; J-F. Blavier, G. Toon, & D. Bekker (2007). An FPGA/SoC Approach to On-Board Data Processing Enabling New Mars Science with Smart Payloads, *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2007.
- Pingree, P.; M. Janssen, J. Oswald, S. Brown, J. Chen, K. Hurst, A. Kitiyakara, F. Maiwald, & S. Smith (2008-a). Microwave Radiometers from 0.6 to 22 GHz for Juno, A Polar Orbiter around Jupiter, *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2008.
- Pingree, P.; L. Scharenbroich, T. Werne, & C. Hartzell (2008-b). Implementing Legacy-C Algorithms in FPGA Co-Processors for Performance Accelerated Smart Payloads, *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, March 2008.
- Platt, J. (1998). Sequential Minimal Optimization: A Fast Algorithm for Training Support Vector Machines, *Microsoft Research Technical Report MSR-TR-98-14*, (1998).
- Sander, S. P.; R. Beer, J-F. Blavier, K. Bowman, A. Eldering, D. Rider, G. Toon, W. Traub, & J. Worden (2008). Panchromatic Fourier Transform Spectrometer (PanFTS) for the Geostationary Coastal and Air Pollution Events (GEO-CAPE) Mission, American Geophysical Union, Fall Meeting, San Francisco, CA, December 2008.
- Yu, H., Y.J. Kaufman, M. Chin, G. Feingold, L.A. Remer, T.L. Anderson, Y. Balkanski, N. Bellouin, O. Boucher, S. Christopher, P. DeCola, R. Kahn, D. Koch, N. Loeb, M. S. Reddy, M. Schulz, T. Takemura, & M. Zhou (2006). A review of measurement-based assessments of the aerosol direct radiative effect and forcing. *Atmos. Chem. Phys.* 6, 613-666.



Aerospace Technologies Advancements

Edited by Thawar T. Arif

ISBN 978-953-7619-96-1

Hard cover, 492 pages

Publisher InTech

Published online 01, January, 2010

Published in print edition January, 2010

Space technology has become increasingly important after the great development and rapid progress in information and communication technology as well as the technology of space exploration. This book deals with the latest and most prominent research in space technology. The first part of the book (first six chapters) deals with the algorithms and software used in information processing, communications and control of spacecrafts. The second part (chapters 7 to 10) deals with the latest research on the space structures. The third part (chapters 11 to 14) deals with some of the latest applications in space. The fourth part (chapters 15 and 16) deals with small satellite technologies. The fifth part (chapters 17 to 20) deals with some of the latest applications in the field of aircrafts. The sixth part (chapters 21 to 25) outlines some recent research efforts in different subjects.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Paula J. Pingree (2010). Advancing NASA's On-Board Processing Capabilities with Reconfigurable FPGA Technologies, Aerospace Technologies Advancements, Thawar T. Arif (Ed.), ISBN: 978-953-7619-96-1, InTech, Available from: <http://www.intechopen.com/books/aerospace-technologies-advancements/advancing-nasa-s-on-board-processing-capabilities-with-reconfigurable-fpga-technologies>

INTECH

open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2010 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](#), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.