The COVE Payload – A Reconfigurable FPGA-Based Processor for CubeSats

Dmitriy L. Bekker, Paula J. Pingree, Thomas A. Werne, Thor O. Wilson, Brian R. Franklin
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109; 818-354-3337
Dmitriy L. Bekker@jpl.nasa.gov

ABSTRACT

The Xilinx Virtex-5QV (V5QV) FPGA is targeted as the rad-hard processor for the NASA Decadal Survey Aerosol-Cloud-Ecosystem (ACE) mission's Multiangle SpectroPolarimetric Imager (MSPI) instrument, currently under development at JPL. With funding from NASA's Earth Science Technology Office (ESTO), JPL is demonstrating a data processing algorithm that reduces the output data rate from MSPI's 9 multi-angle cameras by more than 200x. This real-time on-board processing (OBP) design, based on least-squares fitting, is implemented on the Virtex-5 FPGA. Another ESTO-funded task called COVE (CubeSat On-board processing Validation Experiment) integrates the MSPI OBP algorithm on the V5QV FPGA as a JPL-provided payload to the University of Michigan's M-Cubed CubeSat. One goal of the M-Cubed/COVE mission is to gain on-orbit validation of MSPI OBP on the V5QV FPGA thereby advancing the Technology Readiness Level (TRL) for the ACE mission. M-Cubed/COVE has been selected for launch on the NPOESS Preparatory Project (NPP) mission on October 25, 2011. This paper presents the COVE payload architecture on the M-Cubed CubeSat and explores how the capabilities of this Virtex-5 FPGA-based processing platform could enable future CubeSat missions.

1. INTRODUCTION

The COVE board is based on the Virtex-5QV (V5QV) FPGA and is a secondary payload on the University of Michigan's M-Cubed CubeSat. COVE runs the MSPI on-board processing (OBP) algorithm on image data captured by M-Cubed with its primary camera payload. This section provides a brief overview of MSPI OBP, the Xilinx V5QV FPGA, and M-Cubed.

1.1 Multiangle SpectroPolarimetric Imager

The Multiangle SpectroPolarimetric Imager (MSPI) is an instrument concept in development at JPL to produce a highly accurate multiangle-multiwavelength polarimeter to measure cloud and aerosol properties as called for by the Aerosol-Cloud-Ecosystem (ACE) tier-2 mission concept in the Decadal Survey. This instrument has proposed 9 cameras (8-fixed and 1-gimballed), each of which must eventually process a raw video signal rate around 95 Mbytes/sec over 16-20 channels. A computationally intensive linear least-squares algorithm must be applied to perform data reduction for video processing of the signal output from the photo-detector array. This data reduction can be performed without sacrificing the information content of the camera product for science. The result of the on-board processing algorithm, implemented on the Virtex-5 FPGA, is the reduction of dozens of samples acquired during a 40-msec frame to five parameters. Averaging cross-track and along-track pixels to reduce spatial resolution achieves further data reduction. This technology development is required to enable the process to occur in real-time, with the speed necessary to keep up with MSPI data throughput. For a complete mathematical description of the data processing algorithm, see the journal articles in *Applied Optics*.[1, 2] A 2010 IEEE Aerospace Conference paper provides a concise overview of the algorithm with an emphasis on implementation details.[3]

1.2 Xilinx Virtex-5QV FPGA

The Virtex-5QV is Xilinx's first rad-hard FPGA designed specifically for space applications. It uses different silicon from the commercial and defense-grade Virtex-5 FPGAs, providing protection from single-event upsets (SEU), single-event latchups (SEL), single-event transients, and high total ionizing dose (TID). With these features, triple modular redundancy (TMR) of the entire design is no longer necessary, freeing more FPGA resources for algorithm implementation. As an SRAM-based device, this FPGA is reprogrammable and offers highly desirable processing elements such as digital signal processor (DSP) blocks, high speed clocking resources, and embedded block RAM (BRAM).[4]

Although it is not anticipated that the M-Cubed satellite will go through areas of high radiation, the insertion of the V5QV FPGA into a space-flight mission is an

important technological milestone that will enable its use for future JPL missions, such as on the MSPI instrument for the upcoming ACE mission.

1.3 University of Michigan M-Cubed CubeSat

The University of Michigan's M-Cubed is a 1U[i] imaging CubeSat with a 2 Megapixel CMOS camera to take high-resolution images of the Earth. The satellite will be launched into a low Earth orbit and use a passive magnetic control system to align with the Earth's magnetic field. Once in this configuration, the command and data handling (C&DH) subsystem's Stamp9G20 processor will command the camera to take a picture of the Earth, transfer the image to the COVE payload for on-board processing, and later package the image and OBP results for transmission to the ground station in Ann Arbor, MI. A more detailed concept of operations description is provided in Section 3.2. A U. Michigan in-house developed Electrical Power System (EPS) will power the camera payload, JPL COVE payload board, C&DH, and off-the-shelf radios. The EPS uses onboard solar cells to charge a lithium-ion battery pack through Direct Energy Transfer (DET). All of these subsystems will be housed and supported by a custom-built structure conforming to the CalPoly 1U CubeSat specifications (See Figure 1).[5]

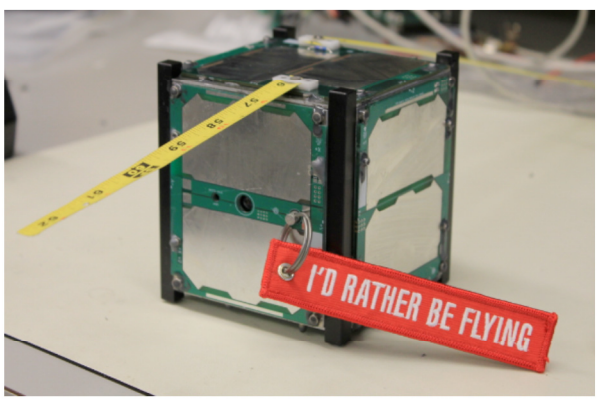

Figure 1. The University of Michigan's M-Cubed CubeSat (Engineering Model), 10 x 10 x 10 cm (Image courtesy of U. Michigan)

[i] CubeSats are characterized by their size and are made in increments of Units, where one Unit, or "1U," is 10x10x10 cm with a nominal mass limit of 1 kg.

2 COVE KEY REQUIREMENTS AND INTERFACE SPECIFICATIONS

As 1U CubeSats are launched in groups of 3 from a common launch system known as the P-POD (Pico-satellite Orbital Deployer), they need to conform to strict requirements in mass and volume. Each 1U of a P-POD deployed CubeSat typically can't exceed 1 kg and must fit within a frame not exceeding 10x10x10 cm. These requirements also imply that CubeSats must utilize power saving design techniques in order to fit within a tight power budget. As a payload on the M-Cubed CubeSat, the COVE board must also adhere to the above specifications.

This section describes the key requirements and interface specifications that drive the design of the COVE payload and its accommodation in the M-Cubed CubeSat and provides an update to information presented at the 2011 IEEE Aerospace Conference.[6]

2.1 Mass

The 1U CubeSat mass specification is typically 1 kg, although M-Cubed has been given an allocation of 1.5kg for the ELaNa[ii] 3 launch. With this overall mass constraint in mind, the M-Cubed design team levied a requirement that the total mass of the COVE payload not exceed 100 grams. The first version of the COVE board, the EM (Engineering Model), has a mass of 99.2 grams. The next version, the FM (Flight Model) has a mass of 83.5 grams (before staking and conformal coat). There were changes in the board design between the two versions, described in more detail in Sections 3.4.

2.2 Power

M-Cubed is able to supply 15W continuously to the entire satellite without depleting the battery. Any power consumption above 15W begins to discharge the battery and should be avoided. The COVE payload running the MSPI OBP algorithm is estimated to consume 3-4W. While this is a significant fraction of the total nominal power available to the CubeSat, the duty cycle of operation for COVE is very short. The COVE board is designed to allow for image data to be transferred to shared on-board memory without turning on the FPGA (the most power-consuming device on the board). The FPGA is only powered on when necessary to process already acquired image data; however, a direct data transfer mode does exist for circumstances when it is desirable to extend the FPGA's operation in the space environment (after primary mission objectives have been met). Nominally, the FPGA will be powered on

[ii] Educational Launch of Nanosatellites (ELaNa)

only long enough to process the image data and transfer its results.

2.3 Board Dimensions

In order to fit comfortably inside the M-Cubed satellite frame, the COVE board was designed to measure 90 x 90 mm. The footprint of the FPGA alone is 42.5 x 42.5 mm. The remaining board space had to accommodate memory, configuration PROMs (programmable read-only memory), regulators, ADCs (analog-to-digital converters) and other support components. The board dimensions and remaining available board space factored heavily in our component selection, described in more detail in Section 3.3.

2.4 Power Interface

M-Cubed provides raw battery voltage in the range of 6.0V-8.4V to the COVE payload. The COVE board regulates secondary voltages required for all of its components, as indicated below:

- FPGA: 3.3V, 2.5V, 1.0V
- Configuration PROM: 3.3V, 1.8V
- Magnetoresistive RAM: 3.3V
- Flash memory: 3.3V
- ADCs: 3.3V, 1.25V
- Buffers, muxes: 3.3V

M-Cubed additionally made available to COVE 3.3V and 5V regulated power with current limits of 2A on each line. While the COVE design for MSPI OBP could work with this dual input power supply interface (as demonstrated in the EM design), there were two reasons why we elected to use the single unregulated power bus input for the final FM design. The first reason is that by not limiting the input power to 2A, future FPGA designs that consume more power than the MSPI OBP design can be accommodated (via upload reconfiguration). Secondly, also to accommodate future COVE payload applications, interface requirements are eased and testing is simplified by requiring only a single supply input of a wide voltage range (we selected dc-dc power converters that can run off any input voltage in the range of 5.5V to 26.5V).

2.5 Communication Interface

COVE communicates with M-Cubed over a 26-pin ribbon cable that is shared with other CubeSat subsystems. Out of a total of 26 I/O lines, 15 are allocated for communication with COVE.

The COVE design is described in Section 3 with Figure 2 showing the I/O signals between the Stamp9G20 and the COVE board (Table 1). Section 3.2 describes the concept of operations that utilizes many of these signals.

Table 1. COVE / M-Cubed Interface

Signal	Function
COVE_VIF_EN	Powers the board interface for image data transfer
COVE_VCC_EN	Powers the FPGA for OBP
SCK, SS, MOSI, MISO[iii] (4)	Serial peripheral interface (SPI) signals; communication interface to Stamp9G20 processor for transferring M-Cubed images to COVE and FPGA-based OBP results to M-Cubed
COVE_HOLD_n	Operational signal that pauses serial connection to the SPI flash device without deselecting it
COVE_CONFIG	Sets FPGA boot mode
COVE_MODE	Sets the SPI transfer mode
COVE_CONFIG_DONE_n	Indicates successful FPGA boot
COVE_PROCESS_DONE	Indicatea successful OBP run
UART_RX/TX (2)	M-Cubed UART to COVE
COVE_DBG (2)	External test access interface

3 COVE DESCRIPTION

The COVE design went through two iterations – Engineering Model (EM) and Flight Model (FM). The sections below detail the FM architecture, concept of operations, component selection, and fabrication and assembly. Differences between the EM and FM versions are also discussed.

3.1 COVE Payload Architecture

The COVE payload architecture, presented in Figure 2, can be subdivided into the following logical subsystems:

- FPGA, PROMs, clock resources
- Shared memory
- Non-shared memory
- System health measurement
- Power
- Isolation buffers

FPGA, PROMS, Clock Resources

Unlike the commercial-grade Virtex-5 FPGA, the V5QV does not have an embedded PowerPC processor. Our MSPI OBP algorithm requires the use of a processor which supervises and controls the algorithm running in the FPGA fabric. On the V5QV, we utilize the MicroBlaze soft-core processor to replace the PowerPC.

Since the V5QV is an SRAM-based volatile FPGA, it requires non-volatile configuration PROMs in order to load a bitstream on power-on. We use two extended

[iii] Clock (SCK), Slave Select (SS), Master Output/Slave Input (MOSI), Master Input/Slave Output (MISO)

temperature range Xilinx XQF32P PROMs to store one FPGA configuration bitstream. These PROMs are re-programmable via a JTAG interface, accessible on the ground only. Before final integration with M-Cubed, we program these PROMs with our latest FPGA firmware and MicroBlaze software.

Only a single clock source is required on the COVE board. We use a 100 MHz oscillator to provide the main system clock for the FPGA. Internally in the FPGA, phase lock loop (PLL) and digital clock manager (DCM) blocks perform all the necessary clock distribution.

Shared Memory

The primary method of transferring image data to the FPGA and processed results to M-Cubed is via a shared flash memory device. A secondary data transfer mechanism exists via an SPI bus transfer directly from M-Cubed to FPGA. In this mode, the FPGA must stay powered on for the data transfer (and subsequent image processing). The Numonyx P5QPCM flash is accessible via an SPI bus. On the COVE board, the flash SPI bus is multiplexed between the FPGA and M-Cubed, allowing either device to have independent access to the shared memory. When accessed by M-Cubed, the FPGA can be powered off.

Non-shared Memory

The COVE board is also populated with non-volatile Magnetoresistive Random Access Memory (MRAM), intended for use as additional instruction and data storage for the MicroBlaze processor and FPGA. Although the OBP algorithm fits entirely within on-chip FPGA memory (BRAM), having additional external RAM enables future memory-demanding designs to be uploaded to the COVE board.

System Health Measurement

Two Analog Devices AD7714 ADCs and a Maxim MAX6627 temperature sensor measure board voltage, current, and operating temperature of the FPGA. All three devices operate over independent SPI buses tied directly to the FPGA. Data from the ADCs and the temperature sensor is recorded by the FPGA and reported to M-Cubed along with the image processing results.

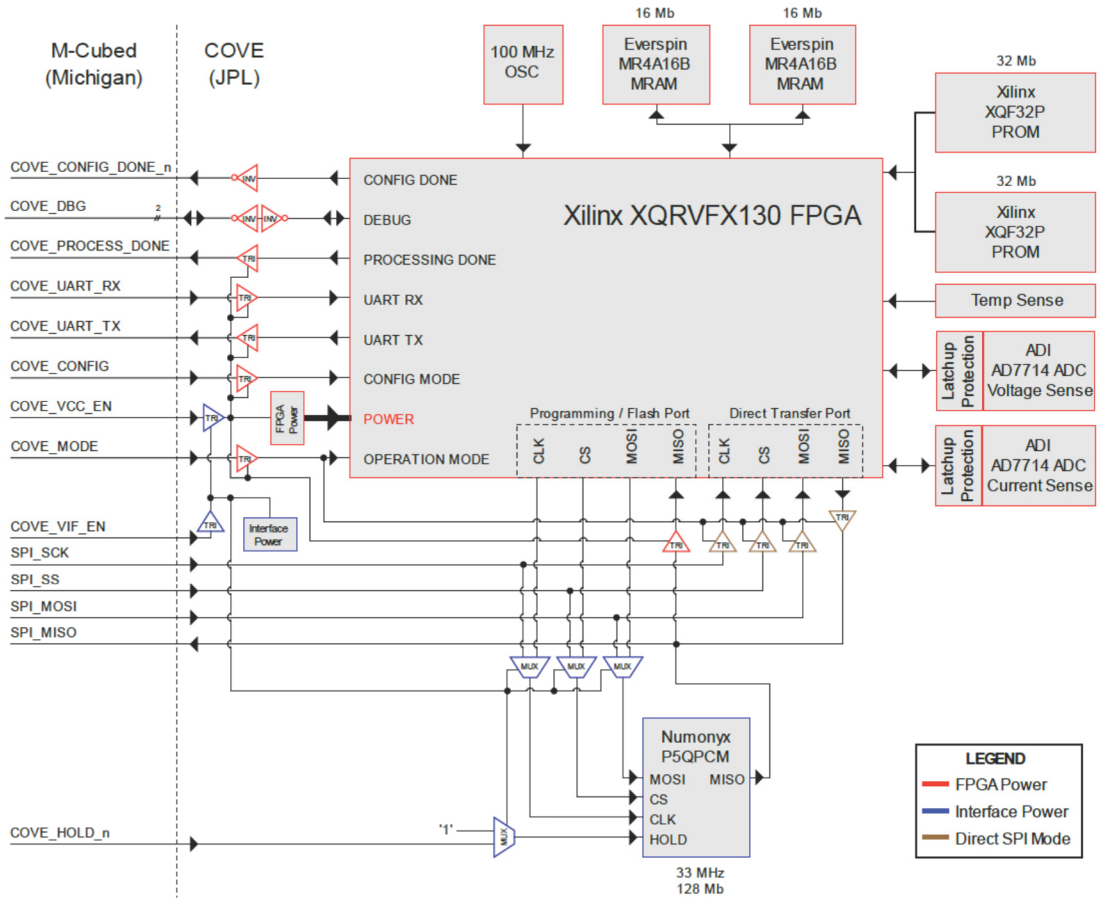

Figure 2. COVE Payload Block Diagram

Power

Two Linear Technology LTM4619 dual 4-amp DC/DC converters are supplied by the unregulated M-Cubed battery voltage and output 3.3V (V_{IF}), 3.3V (V_{CC}), 2.5V, and 1.0V secondary voltages. A single Linear Technology LT3029 dual low-dropout (LDO) regulator generates the remaining low-current secondary voltages by first regulating the 2.5V rail down to 1.8V and then 1.8V down to 1.25V. Enable pins on these parts are used to switch power on/off eliminating the bulky MOSFET based switching circuits used on the EM board. Secondary supply currents are monitored with Maxim MAX9634 current-sense amplifiers driving the ADC.

Isolation Buffers

Since the FPGA on the COVE board is not powered on at all times, special care must be taken not to exercise I/O lines tied to the FPGA while it is in the off state. We utilize tri-state buffers and inverters to isolate the FPGA from M-Cubed and other board components that may be in a powered state when the FPGA is off. Additionally, these buffers are capable of the bus-hold[iv] functionality, ensuring that all I/O to other devices on the COVE board is held at a non-floating logic level.

3.2 COVE Concept of Operations

The COVE board is designed with the intent to minimize power consumption while offering the full data processing capability of the V5QV FPGA. This is achieved by selectively powering on/off the interface and the FPGA. In the section below, parentheses indicate the specific component in the M-Cubed/COVE design, but can be generalized for other CubeSat applications. Refer again to Figure 2 for component and signal references.

COVE Primary Mission

During most of M-Cubed operation, the COVE board is powered off. Although the voltage regulators receive power from the battery, their control lines are set to disable voltage output.

When M-Cubed has acquired an image (taken with its on-board OmniVision camera) and is ready for on-board processing, a command is sent to COVE to turn on its SPI interface and flash memory. M-Cubed then writes the raw image data to a predetermined address in the shared flash memory. This process can take nearly 3 minutes for a standard 5.76 MB image because the (Stamp9G20) microprocessor on M-Cubed is limited in its SPI transfer rates to approximately 30 seconds per MB. The power consumption in this step, however, is very low (~0.25W) as only a few devices on the COVE board are in the on state.

Once the full image has been transferred to the shared flash memory, the (Stamp9G20) microprocessor gives up control of the shared flash memory and commands COVE to turn on the FPGA. The V5QV FPGA loads the configuration bitstream from the PROMs and begins to read and process the data from the shared flash memory. The results of the (MSPI algorithm) data processing are written back to the shared flash memory into a predefined address space.

The FPGA notifies the (Stamp9G20) microcontroller when it is done processing data. The microcontroller then acquires control of the shared flash memory and commands the FPGA to power off. The processed data is then read back by the (Stamp9G20) microcontroller and the COVE payload activity is complete. As the final step, the microcontroller commands the COVE payload to power off its SPI interface and flash memory.

COVE Future Use

Once the primary mission is complete, the COVE board could be made available to other customers to demonstrate their FPGA-based technology on the V5QV FPGA. To accommodate future uses of the COVE payload, the board is designed with the ability to upload new configurations to the FPGA while M-Cubed is on orbit. In addition to the XQF32P PROMs that store the "golden" bitstream (non-updateable on orbit), the FPGA is also able to configure from the shared flash memory. A region of the shared flash memory is reserved for storing a secondary configuration bitstream. This bitstream may be uploaded to the flash memory from the microcontroller in the same manner image data is transferred to this memory.

COVE also provides non-volatile MRAMs for external code and data storage as may be required by future memory intensive MicroBlaze applications. This use case is enabled by first uploading a new configuration bitstream into the shared flash memory, including additional content for insertion into the MRAMs. Upon configuration, the FPGA can be instructed to copy certain contents of the shared flash memory into the MRAMs and then commence MicroBlaze processor boot.

[iv] For more information on bus-hold, refer to: http://focus.ti.com/lit/an/scla015/scla015.pdf

3.3 *Component Selection*

Component selection was a very important process during the COVE design phase. The main objectives in choosing appropriate components were: feature set, package options, ease of integration, and extreme environment operation. While the rad-hard V5QV FPGA is the most robust component on the COVE board, we wanted to make sure other commercial and military-grade components were well suited for the mission:

- ***SPI flash memory:*** the selected component is a phase change memory (PCM) device known from previous designs to work with Xilinx FPGAs. Flash devices trap electrons to store information; therefore, they are susceptible to data corruption from radiation. However PCM exhibits higher resistance to radiation so this is an advantage for space applications.[7] In addition, the SPI flash supports bit-alterable writes, which is an efficient flash memory writing technique in newer devices. This mode enables writing data to flash at higher rates.

- ***Muxes and tri-states:*** the selected components exhibit a measured propagation delay of 4-5 ns. A low propagation delay is necessary to support fast flash memory access. Furthermore, the selected tri-state buffers feature "bus-hold" circuitry which reduces the use of external pullup/down resistors.

- ***Xilinx PROMs:*** these components have been used in previous designs. A version was selected with an extended temperature range.

- ***ADCs:*** selected for their SPI interface to the FPGA, low pin count, small package size, and low power consumption.

- ***MRAMs:*** selected for their small package size, non-volatile storage, better radiation tolerance, and previous space flight.[8]

- ***Power circuitry:*** provided by two ruggedized, dual output DC-DC converters for the 3.3V, 2.5V and 1.0V secondary voltages, and two low current, linear regulators for the 1.8V and 1.25 secondary voltages.

3.4 *COVE EM/FM Differences*

The COVE payload went through one re-design after the initial EM was built. Initial design and component selection dictated that all required components could not fit on a single board. The EM design segregated all power regulators on a separate smaller daughter board that was mounted on stand-offs above the full-sized FPGA-based processor board (Figure 3). Furthermore, only one ADC was used on the EM board to measure a smaller set of voltages and currents.

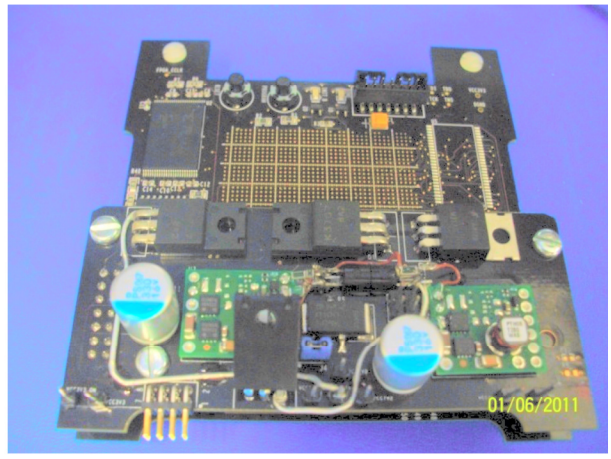

Figure 3. EM COVE Payload
(regulator board on top, processor board below)

In the FM design, we selected smaller MRAMs and power regulators, and performed an overall optimization of the layout. Doing so, we were able to fit all components on a single board and include an additional ADC for more voltage and current monitoring. Additional buffers were also added to provide isolation on the M-Cubed interface I/O. The mass of the board decreased after going to a single board solution.

With two separate printed circuit boards (PCB), the revised COVE board design accommodates either the commercial-grade, ball grid array (BGA) Virtex-5 FPGA or the space-grade, column-grid array (CGA) V5QV device.

3.5 *Flight Model (FM) Assembly*

Figure 4 shows the assembled Flight Model COVE payload board with an Engineering Sample (ES) of the V5QV device. As shown, a transistor is taped to the center of the FPGA and hay-wired to the board circuitry to allow temperature sensing of the device. This is required only for the ES part. With the V5QV production part, temperature sensing is handled internal to the device.

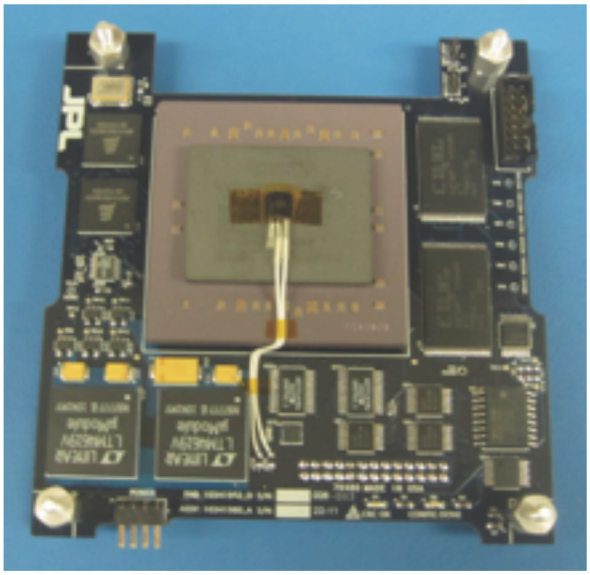

Figure 4. FM COVE Payload (single board); transistor haywired to the top of the FPGA provides external temperature sensing.

The COVE payload board contains two flight headers – one for power and one for command/data – to interface to the M-Cubed EPS & C&DH subsystems, respectfully. A third header provides a JTAG test interface directly to the FPGA. Structural mounting stand-offs are seen in each corner providing the mechanical interface to M-Cubed. The large notch seen on the top of the board is a mechanical keep-out zone where the M-Cubed camera will be mounted inside the assembled CubeSat. Future versions of the COVE board where such a keep-out zone is not required will allow for additional PCB area to accommodate other components or headers, as needed.

4 TESTING

The COVE payload is tested with the help of an off-the-shelf Digilent Spartan-3 Development Board (Figure 5). The Spartan-3 board emulates the M-Cubed I/O interface, allowing us to test the board without having access to the complete CubeSat assembly. All I/O lines can be exercised to verify the full command set, and the SPI flash memory can be independently read/written for image and configuration bitstream transfers. The COVE UART and debug lines are routed through the Spartan-3 board for user access.

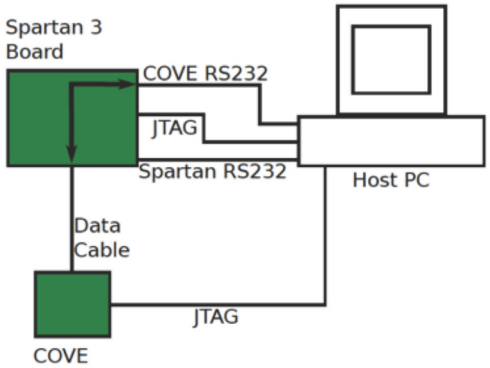

Figure 5. COVE Lab Test Configuration

5 LAUNCH SELECTION

The CubeSat On-board processing Validation Experiment (COVE), integrated on the M-Cubed CubeSat, has been selected as part of the ELaNa 3 manifest for launch as a secondary payload on the NPOESS Preparatory Project (NPP) Mission. This selection was part of the NASA Space Operations Mission Directorate's (SOMD) 2010 CubeSat Launch Initiative. The NPP satellite will be launched from the Western Range at Vandenberg Air Force Base from SLC-2, California, by a Boeing Delta II-7920-10 launch vehicle. The NPP mission will be launched into a high inclination, sun-synchronous orbit. NPP is currently schedule for launch on October 25, 2011.[9]

6 CONCLUSION

JPL has been working on a Multiangle SpectroPolarimetric Imager (MSPI) instrument to fulfill hardware needs for the Aerosol-Cloud-Ecosystem (ACE) Decadal Survey mission. In order to meet the real-time data processing requirements, JPL has developed an algorithm for MSPI that will run on the Xilinx Virtex-5 FPGA. The Virtex-5QV (V5QV) FPGA is not yet space-flight tested; however, this in-flight validation of the technology on a pre-cursor CubeSat mission is extremely valuable toward advancing the technology readiness levels of both the V5QV device and the polarimetric OBP algorithm for MSPI and the ACE mission. To this end, JPL and the University of Michigan are collaborating to use the M-Cubed CubeSat as an on-board processing technology validation platform. M-Cubed, with the JPL COVE payload, will be launched into low Earth orbit as a secondary payload of the NPP mission.

MSPI and other future JPL science instruments that will rely on the hardware acceleration provided by FPGA embedded technology to meet their on-board processing

requirements can benefit from the low-cost capability of CubeSats to carry small payloads into LEO.

ACKNOWLEDGMENTS

The authors acknowledge the contributions of the following Space Grant summer students who helped advance the prototype design of the COVE payload board during their 2010 internship at JPL: Brad Freyberg (University of Michigan), Vikram Ivatury (University of Michigan), and Jesse Frey (University of Alaska, Fairbanks).

The authors acknowledge the collaborative contributions of the University of Michigan M-Cubed CubeSat team, led by Professor Jamie Cutler and student Michael Heywood. Dr. Cutler and Mr. Heywood are Co-Investigators of the COVE task.

The authors acknowledge the support of the JPL MSPI team lead by David Diner, Sven Geier and previously, Ab Davis.

The authors acknowledge the support of NASA's Earth Science Technology Office (ESTO), including Mike Pasciuto and Charles Norton.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration, and by the University of Michigan.

©2011 California Institute of Technology. Government sponsorship acknowledged.

REFERENCES

1. Diner, D. J., A. Davis, B. Hancock, S. Geier, B. Rheingans, V. Jovanovic, M. Bull, D. Rider, R. A. Chipman, A-B Mahler, and S. McClain, "First results from a dual photoelastic modulator-based polarimetric camera," *Appl. Opt.* 49, 2929-2946, 2010

2. Diner, D. J., A. Davis, B. Hancock, G. Gutt, R. A. Chipman, and B. Cairns, "Dual photoelastic modulator-based polarimetric imaging concept for aerosol remote sensing," *Appl. Opt.* 46, 8428-8445, 2007

3. Werne, T. A., D. L. Bekker, and P. J. Pingree, "Real-Time Data Processing for an Advanced Imaging System Using the Xilinx Virtex-5 FPGA," *Aerospace Conference, 2010 IEEE,* March 2010

4. Virtex-5QV FPGA Product Brief, Xilinx, Inc., http://www.xilinx.com/publications/prod_mktg/virtex5qv-product-brief.pdf

5. Heidt, H., J. Puig-Suari, A. S. Moore, S. Nakasuka, and R. J. Twiggs, "CubeSat: A New Generation of Picosatellite," in *Proceedings of the 15th Annual AIAA/USU Conference on Small Satellites*, Logan, UT, August 2001

6. Pingree, P., D. Bekker, T. Werne, T. Wilson, "The Prototype Development Phase of the CubeSat On-board Processing Validation Experiment (COVE)," *Aerospace Conference, 2011 IEEE,* March 2011

7. Phase-change memory, Wikipedia, http://en.wikipedia.org/wiki/Phase_change_memory

8. Angstrom Aerospace uses Everspin MRAM, http://www.everspin.com/technology.php?qtype=11

9. ESTO Technology Selected for Flight Validation on CubeSat for Launch on NPOESS Preparatory Project Mission, in *ESTO Weekly*, August 20, 2010