

# Development of the Malleable Signal Processor (MSP) for the Roadrunner On-Board Processing Experiment (ROPE) on the Tacsat-2 Spacecraft

R.L. Coxe<sup>1</sup>, G.H. Romero<sup>1</sup>, A. Pakyari<sup>1</sup>, M. Leary<sup>2</sup>, J. Lyke<sup>3</sup>, and D. Fronterhouse<sup>4</sup>

<sup>1</sup>Physical Sciences Inc.

<sup>2</sup>Newgrange Design Inc.

<sup>3</sup>AFRL/VSSE

<sup>4</sup>Scientific Simulations Inc.

## Abstract

In this paper, we describe the design and development effort at Physical Sciences Inc. (PSI) of the Malleable Signal Processor (MSP), a reconfigurable computing engine equipped with five radiation-tolerant Xilinx XQR2V3000 FPGAs slated for the Roadrunner On-Board Processing Experiment (ROPE) imaging payload onboard the AFRL TacSat-2 satellite. TacSat-2, slated for launch in 2006, is one of the first flagship missions of the DoD Office of Force Transformation's Responsive Space Initiative. We describe the mission objectives of TacSat-2 and provide an overview of ROPE and MSP functionality. We conclude with lessons learned while developing new technology in the Brave New World of Responsive Space.

## 1. Introduction

In 2003 the Air Force Research Laboratory Space Vehicles Directorate identified PSI's FPGA reconfigurable computing design concept, originally formulated in a Phase I Small Business Innovation Research (SBIR) program, as a promising candidate to meet the objective of the Roadrunner On-board Processing Experiment (ROPE). ROPE is intended to demonstrate on-board, real-time, adaptive processing of multispectral imagery. The MSP supports on-the-fly, pipelined radiometric calibration, JPEG image compression, and anomaly detection on 4 bands of multispectral imagery data (12 bits/pixel, 6144 pixels/line, input pixel clock rate of ~60 MHz) via 4 parallel pipelines in real time. Additional capabilities of the MSP include rapid prototyping and on-demand functional upgrades on-orbit.

The AFRL TacSat-2 mission [1], slated for launch in 2006 on the SpaceX Falcon launch vehicle, is taking place under the aegis of the DoD Office of Force Transformation's Responsive Space Initiative [2]. The TacSat missions are intended to demonstrate the feasibility of the goals of the Air Force Responsive Space Initiative: develop spacecraft in 6-12 month time frames, keep spacecraft in inventory for less than 3 years, transition them from stored state to orbit in one week or less, perform autonomous on-orbit checkout in less than 24 hours, and employ modular design methodologies and standard interfaces. In the context of the "faster-better-cheaper" triad, the TacSat objectives are geared towards optimizing missions in the "faster-cheaper" part of phase space. Responsive support to tactical users is also a key goal of Responsive Space-- capabilities such as field tasking, collection, and downlink of mission data in a single pass and dynamic re-tasking based on both external commands and autonomous on-board cueing.

In this paper, we describe the salient features of PSI's MSP system, our system development methodology, and the technical and programmatic challenges we faced in the severely cost- and schedule-constrained environment of one of the first Responsive Space technology demonstration programs.

## 2. Malleable Signal Processor (MSP) System Overview

### 2.1 Roadrunner On-Board Processing Experiment

The Roadrunner On-Board Processing Experiment (ROPE) is a reconfigurable on-board, real-time multispectral image processing system with a throughput of 360 megabytes per second. As pictured in Figure 1, ROPE consists of 4 major components: a wide-field multispectral imager, with red, green, blue, and panchromatic bands; the MSP, the Fusion Processor, or FP, and an 8 GB solid state buffer—a large array of SDRAMs. The imager was built by Nova Sensors in Solvang, CA. Space Micro Inc. of San Diego, CA is supplying the FP and the solid state buffer. The MSP is built around five radiation-tolerant Xilinx Virtex-II FPGAs [3]. The MSP has 4 processing pipelines, one for each imagery band. A Xilinx MicroBlaze [4] 32-bit embedded RISC processor in Service FPGA acts as the master MSP controller. The first processing step is a multiplicative gain correction and a subtractive offset correction to compensate for non-uniformities in the focal plane array. The MSP then performs lossy JPEG image compression [5]. In the future, the MSP will also be capable of executing an RX anomaly detection algorithm [6,7] prior to compression. The compressed data, as well as pixels flagged by the RX algorithm, are transferred to the FP via a 64-bit parallel interface. Alternatively, data can be packetized and serialized and directly downlinked to the Common Data Link (CDL) RF modem. Raw data can be buffered in SDRAM on the MSP and transferred to the FP—buffering is necessary because the clock provided to the MSP from the FP across the 64-bit interface is not fast enough to keep up with the MSP raw data flow. When the data is compressed by a factor of  $>4:1$ , however, buffering is not necessary. Data transferred to the FP is subsequently stored in the solid state buffer. Another data path supported by the MSP is to receive data from the FP across the 64-bit interface either buffer it in SDRAM, downlink it to the CDL, or store it as a FAT16 file on a CompactFLASH card.

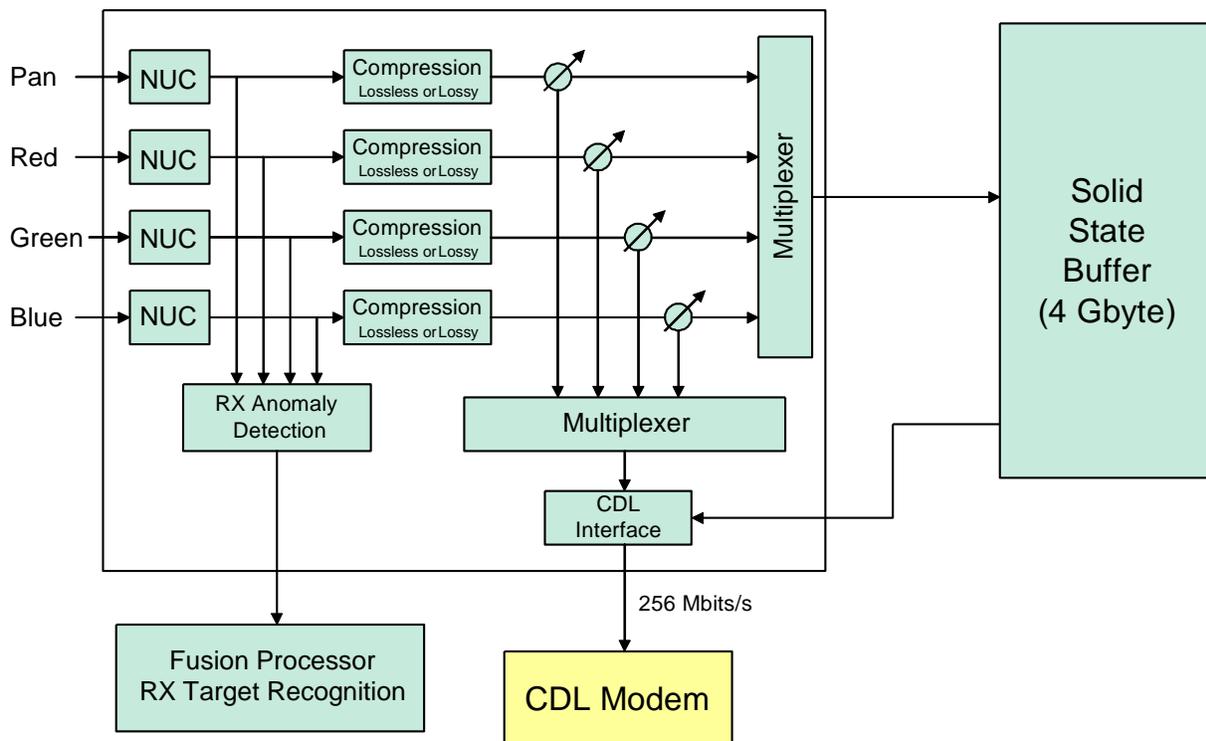


Figure 1. Block Diagram of the Roadrunner On-Board Processing Experiment (ROPE).

## 2.2 Malleable Signal Processor (MSP) Feature Summary

A photograph of the MSP Flight Engineering Model board appears in Figure 2. A block diagram of key MSP components and data paths is shown in Figure 3. Table 1 summarizes the FPGA logic resources available on the MSP. The salient features of the Malleable Signal Processor (MSP) are:

- 5 radiation-tolerant Xilinx XQR2V3000-4BG728N Virtex-II FPGAs
- MicroBlaze 32-bit embedded RISC processor in Service FPGA acts as master MSP controller
- 2 external 512 K x 32 SRAMs for NUC coefficient storage for the 2 front-end FPGAs
- Crosspoint 68-bit interconnect, buses between front & back end FPGAs and Service FPGA
- 32Mx32 external SDRAM for each of the 4 processing elements
- 32Mx16 external SDRAM for Service FPGA
- 100 MHz and 25 MHz clocks
- Real-time clock for time synchronization
- 3.3 V and 1.5 DC power supply
- 32Mx16 FLASH PROM for configuration bitstream storage
- Up to 2 GB CompactFLASH card with SystemACE interface to Virtex-II FPGA
- 64-bit bi-directional data interface to Fusion Processor
- 1 pulse per second timing interface
- 4 x 60.28 MHz image data serial → 12-bit parallel inputs
- 2 RS-232 serial interfaces
- Radiation-tolerant configuration PROMs for redundant Service FPGA power-up
- External Interfaces:
  - Space Micro Inc. Fusion Processor sources commands to MSP via a 230.4 kbps serial port and orchestrates data storage on an 8 GB solid-state buffer via a 125-pin Z-Pack connector.
  - 6 x 42.8 MHz ECL PMED serial data downlink interfaces on the Common Data Link (CDL) RF modem.

Table 1. MSP FPGA Logic Resources

System Gates	~15,000,000
Logic Cells	161,280
BlockRAM (kbits)	8640
18x18 Multipliers	480
Digital Clock Managers	60
Maximum Distributed RAM (kbits)	2240
User I/O Pins	2580

The MSP is 9.5” x 8” and weighs approximately 1 lb without the frame and about 3 lbs with the frame. Each of the five radiation-tolerant Xilinx Virtex-II FPGAs provides ~3 million reconfigurable logic gates. All other MSP components, with the exception of 3 radiation-tolerant Xilinx PROMs that act as a redundant configuration mechanism for the Service FPGA, are military/industrial temperature grade COTS components. The MSP receives serialized data from the imager and deserializes each of the 4 imagery bands. It also serializes and packetizes data to each of the 6 CDL downlink channels. The peak power consumption of the MSP is ~20 W, but it is typically in the 5-15 W range. Because the MSP will be operating in the vacuum of space, the flight hardware has copper strips bonded to the 5 FPGAs

with thermal cement and bolted to the aluminum frame. The MSP has 4 on-board temperature sensors and built-in-self-test functions. The results of the self-test and periodic temperature readouts are stored as a telemetry log file to the CompactFLASH card and periodically transferred to the FP when instructed. The FP supplies the MSP with 6A of 5 V DC power.

The MSP supports a plethora of different data paths: The MSP receives and acknowledges commands from the FP via a 230.4 kilobits-per-second serial port. Non-uniformity correction coefficient files and new MSP personalities can be uploaded from the FP, buffered in the SDRAM attached to the Service FPGA, and stored to the CompactFLASH card. The NUC files are software-selectable and, under command from the FP, the coefficients are uploaded to the SRAMs attached to FPGAs #1 and #2, where the NUC functional modules reside. Raw data is buffered in FPGAs #3 and #4 and can be sent to the CDL interface, the FP, or stored on the Compact FLASH card via the MicroBlaze. Compressed imagery can be either sent to the FP or directly downlinked. The MSP reorders the pixels from the 4 imagery bands following the radiometric calibration, since the readout from the imager is not multiplexed in a straightforward manner.

The MSP has six interfaces to the outside world: the FP command interface, the 64-bit bi-directional parallel data interface to the FP, a 1 PPS timing interface from the imager to the FP, the six CDL serial downlink channels, and the 4 imagery inputs. The ROPE system is set up such that the MSP acts as a slave processor to the FP. The MSP initiates operations only when instructed to do so via the FP command interface. It should be noted that the MSP is fully capable of acting as its own master in alternative system configurations.



Figure 2. Photograph of the MSP Flight Engineering Model.

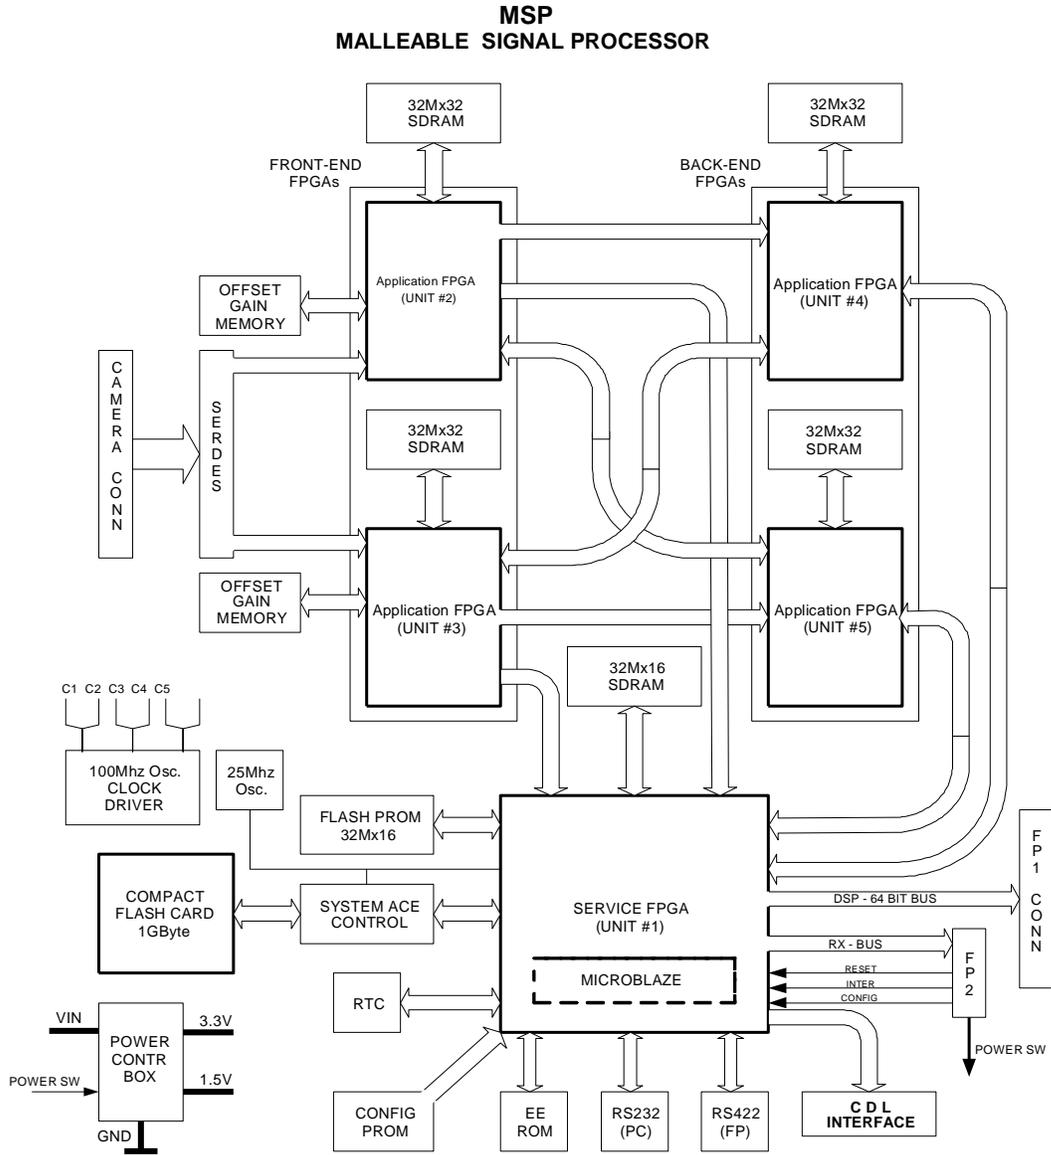


Figure 3. Block diagram of key MSP components and data paths.

### 2.3 MSP Operational Modes

PSI delivered the MSP to the Air Force with 3 primary operational modes, or “personalities”:

- Personality #1: 16:1 Lossy JPEG
- Personality #2 : 4:1 Lossy JPEG
- Personality #3: Calibration
  - No compression core
  - Buffering in SDRAM in FPGAs #3 and #4 of 200 ms of raw data
- Personality #4: 16:1 Lossy JPEG + RX (RX anomaly detection not yet supported)
- Personality #5: ~2:1 Lossless JPEG (compression core developed for future implementation)

The non-uniformity correction takes place for all personalities except for #3. The default MSP mode is lossy JPEG with a compression ratio of ~16:1. JPEG compression with a compression ratio of ~4:1 is also supported. A calibration personality enables the collection of 200 ms of raw imagery data to allow the FP to compute new non-uniformity correction gain and offset coefficients. A full-featured mode, which will be supported in the future, will consist of RX anomaly detection in conjunction with ~16:1 Lossy JPEG. For pixels flagged by the RX algorithm, the calibrated, uncompressed pixel values will also be transferred to and stored by the Solid State Buffer.

Since the FPGAs on the MSP are reconfigurable, it is possible to upload personality updates to the MSP from the ground via the FP and store them on a CompactFLASH card on the MSP board. PSI selected the Xilinx SystemACE CF chip, coupled with a 2 GB CompactFLASH card to orchestrate the configuration of the MSP. The SystemACE chip also enables the MicroBlaze embedded processor in the MSP Service FPGA to access the CompactFLASH card as a storage device.

SEU mitigation is a popular topic at MAPLD, but for the ROPE system, we did not have the time or the resources to implement TMR or bitstream scrubbing. TacSat-2 is a “capabilities-driven” mission: time, manpower, and cost constraints precluded extensive SEU mitigation. Ergo, the strategy on-orbit will be for the FP to reconfigure or power-cycle the MSP in the event the MSP fails to respond to a watchdog timer sent at 100 ms intervals.

## 2.4 MSP Software-Adjustable Parameters

The MSP has many adjustable parameters that can be modified via commands from the FP that are processed by the command parser routine executing on the MicroBlaze. The MicroBlaze orchestrates the configuration of the MSP and the flow of data through it via software I/O to internal status and control registers in each of the 5 FPGAs as well as IBM On-chip Peripheral Bus (OPB) driver routines. For instance, a GPIO VHDL peripheral is attached to the MicroBlaze, which is connected to the EEPROM, the Real-time clock, and the 4 temperature sensors, which enables software configuration and readback of these external devices.

Software-adjustable MSP parameters include:

- # CDL channels enabled (3 or 6)
- # Imager channels to capture (4 or 1)
- Image capture (enabled or disabled)
- Transfer data from Solid State Buffer to MSP for CDL downlink
- NUC Table # (for configuration and uploads)
- MSP Personality # (for configuration and uploads)
- Start Time and Duration of image capture
- RX Threshold
- Enable direct downlink from MSP to CDL
- Enable RX data transfers to Fusion Processor (FP)
- Enable 64-bit data transfers from FP to MicroBlaze data bus and store data as FAT16 binary files on the CompactFLASH card for storage of NUC tables, new personalities, data diagnostics, rapid prototyping/debugging.
- Configure raw data SDRAM buffering via FPGAs #3 & #4 (read & write burst sizes, pipeline delays).
- Select image source (Nova Imager or synthetic image stored in FPGA #1 & #2 SRAM). If synthetic image is selected, the red user LED will flash.

### 3. System Development Methodology and Lessons Learned

#### 3.1 System Integration Issues

Responsive Space is the wave of the future, but since TacSat-2 is one of the first major spacecraft development projects, the process is still a work-in-process. PSI succeeded in implementing significant additional functionality and new data paths to the MSP system that were introduced after the design effort was well underway. Finding the perfect balance between limited manpower, money, and time and increasingly demanding system requirements was and will continue to be a challenge. The major technical hurdles invariably appeared at the interfaces between hardware built by different companies. We all underestimated the engineering labor and system engineering coordination that ended up being necessary to resolve these issues, particularly given the geographical separation of the various participants. The importance of comprehensive, non-ambiguous Interface Control Documents and exhaustive test reports, as well as a decisive, well-organized, proactive lead system integrator cannot be overemphasized. We made extensive use of digital photographs and oscilloscope and logic analyzer screen captures to document system configurations and system behavior.

PSI's decision to fabricate multiple hardware units (10 boards in 2 EM turns), while costlier than just producing 1 or 2 boards, was definitely beneficial to the bi-coastal integration effort. We made a concerted effort to strictly adhere to a standard VHDL file hierarchy and made use of standard templates for functional module development. All VHDL modules were designed with reusability in mind. We did not use a formal configuration control system, but recognize the wisdom of doing so in the future.

#### 3.2 Use of Third-party IP Cores

In what we initially thought would be a time-saving measure, we made extensive use of 3rd party IP cores. Doing so was not as straightforward as we hoped at the outset. Many at MAPLD have mentioned this point in the past, but it is worth revisiting. Assessing the integration complexity of a 3rd party IP core *a priori* is not an easy task, particularly since the provider is often unwilling to share too many details about a core until it is purchased, often at significant cost. PSI incorporated five different varieties of IP cores into the MSP design:

1. The Xilinx COREGEN cores included in the Xilinx Integrated Synthesis Environment (ISE) development tools.
2. The MicroBlaze soft-core 32-bit RISC processor included with the Xilinx Embedded Development Kit (EDK).
3. At the Air Force's behest, we purchased the lossy JPEG core from Amphion Semiconductor (Belfast, Northern Ireland).
4. We commissioned Birger Engineering (Boston, MA) to create a lossless JPEG core (LOCO-I).
5. PSI developed the RX core in-house.

We had no issues whatsoever with the Xilinx COREGEN cores, which tend to be common logic elements: counters, FIFOs, registers, etc. The COREGEN graphical software package is straightforward and enables users to parameterize each core for their specific use. Many cores can be generated as relationally placed macros, which often facilitate the satisfaction of stringent timing requirements and accelerate the place and route process.

The MicroBlaze aspect of the MSP development, however, was particularly challenging due to the novelty (and continually evolving nature of) the development tools. The interfaces between hardware and software in early versions of the tools were quite problematic. In particular, at the outset we faced many problems with clock signal connections common to both the MicroBlaze and its peripherals and the

surrounding custom VHDL logic. We spent countless hours writing diagnostics and creating workarounds. But the release of the Xilinx EDK version 6.3 and ISE version 6.3 tools resulted in vast improvements in reliability and ease-of-use. The ability to instantiate the MicroBlaze core as a sub-module in an ISE project alleviated all of the timing issues. However, we periodically encountered bizarre software behaviors that obviously arose from integrating EDK elements into the ISE software. Trial and error, persistence, and good documentation habits enabled us to eventually master these tools to design embedded applications for Xilinx FPGAs. In retrospect, the ability to simulate the MicroBlaze within the ModelSim VHDL simulator would have been invaluable for the integration and test process. Preparing and executing ModelSim simulations of the MicroBlaze component of a design will become standard operating procedure for future integration and test efforts.

We eventually succeeded in getting the lossy JPEG core to perform as advertised, but spent the better part of 6 months integrating the four cores into the fabric of the MSP processing pipeline. Even with the VHDL source code provided, the integration process was much more involved than simply dropping the core into the design, primarily due to the strict timing requirements in the image processing pipelines necessary to ensure real-time operation. The documentation that came with the core was incomplete and often ambiguous and technical support was inconsistent and often untimely. The quoted timing performance of the core, we discovered, was not based on actual measurements in hardware, rather the output of the synthesis tool. We spent a considerable amount of effort adding pipelining to the core in an effort to increase timing performance to the necessary levels. We finally resorted to increasing the clock frequency of the entire system, which added complexity at the front-end interface with the imager electronics and the back-end interfaces with the Fusion processor and the Common Data Link RF downlinks in order to handle the clock domain crossings properly.

The lossless JPEG core never met timing and was removed from the baseline system design. In the future, we have the option of rewriting it and offering lossless JPEG as an MSP upgrade. The subcontractor was not familiar enough with the inner workings of the MSP system to fully appreciate subtle timing issues. In the future, if PSI commissions an external entity to generate IP for our FPGA designs, we will insist that that development take place on-site, so the subcontractor can interact with the design team and resolve potentially fatal design flaws that can arise from “throwing the design over the wall.”

PSI developed a VHDL implementation of the RX algorithm in-house. Since the logic resources were not available to implement a floating-point version and simultaneously perform image compression, a requirement insisted upon by the Air Force, we were constrained to fixed-point arithmetic. The RX algorithm, like many anomaly detection operations, involves a matrix inversion, a problematic operation in fixed-point math which may not have adequate precision to generate meaningful anomaly detections. Furthermore, in an effort to compress the RX module into the available logic resources, we came to the conclusion that it was impossible to meet the image pipeline timing requirements. The outcome of the RX development is a prime example of the consequences of “requirements creep.” When the MSP board was originally designed, PSI was not informed that the Air Force intended the system to perform real-time anomaly detection. We undertook the RX development in an effort to be responsive to our customer, but in the end spent considerable time, money, and effort going down an unsuccessful path. The attention we devoted to RX diverted already limited resources away from core development tasks. In retrospect, we have learned that it almost never pays off not to trust our gut reactions regarding the feasibility of an add-on technical task, particularly if additional resources are not forthcoming to cover both the development costs and account for the domino effect that occurs during integration following an *ex post facto* insertion into a pre-existing design.

#### 4. Conclusions

In this paper, we described the design and development effort at PSI of the Malleable Signal Processor (MSP), a reconfigurable computing engine equipped with five radiation-tolerant XQR2V3000 FPGAs slated for the Air Force Research Laboratory TacSat-2 imaging payload. AFRL/VSSE identified PSI's FPGA reconfigurable computing design concept, originally formulated in a Phase I Small Business Innovation Research (SBIR) program, as a promising candidate to meet the objective of the Roadrunner On-board Processing Experiment (ROPE). ROPE is intended to demonstrate on-board, real-time, adaptable processing of multispectral imagery. The MSP supports on-the-fly, pipelined radiometric calibration, JPEG image compression, and anomaly detection. Additional capabilities of the MSP include rapid prototyping and on-demand functional upgrades on-orbit.

The AFRL TacSat-2 mission, currently slated for launch in 2006, is a flagship mission of the DoD Office of Force Transformation's Responsive Space Initiative. The stated goal of Responsive Space is to field satellite payloads in weeks and months, rather than decades. Responsive Space represents a shift towards satellite missions that are "faster, cheaper, and good enough." We presented a series of lessons learned while developing new technology in the Brave New World of Responsive Space, many of which came to our attention during the process of overcoming the challenges of a severely resource and schedule constrained bi-coastal system integration effort without centralized system engineering coordination and the incorporation of IP cores into a reconfigurable computer system design.

The MSP hardware platform itself can be readily retargeted to other reconfigurable computing applications without major hardware modifications. PSI is exploring sonar beamforming and other pipelined FFT processing applications, our original goal of real-time HSI processing (a notional scenario is displayed in Figure 4), more sophisticated anomaly and edge detection, and neural computation engines.

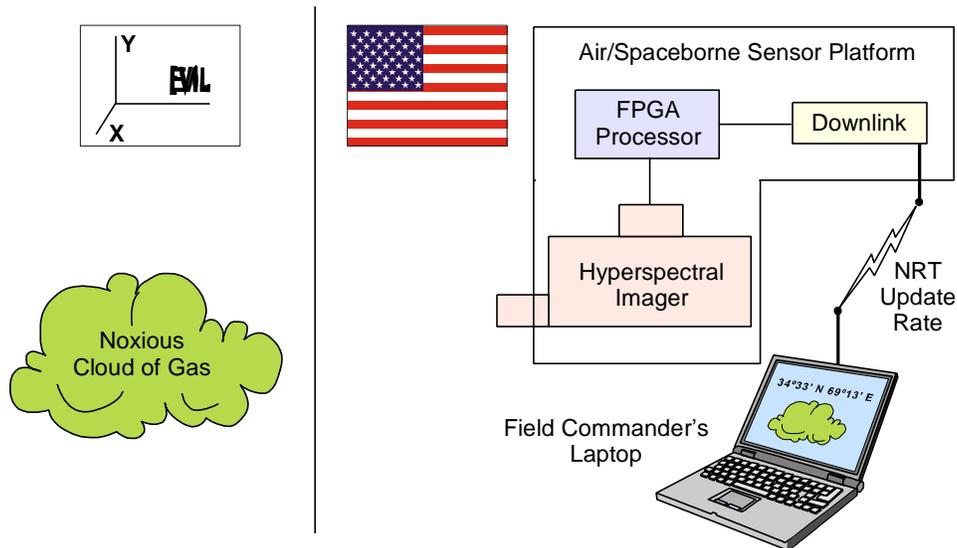


Figure 4. A notional real-time hyperspectral imaging scenario enabled by the MSP.

## Acknowledgements

We could not have succeeded in designing, building, and delivering the MSP flight hardware and firmware in <20 months without Air Force Phase II SBIR and Phase II Enhancement sponsorship (Contract No. F24601-03-C-0086). We would like to recognize Catalyst Manufacturing of Endicott, NY for their high-quality PC board fabrication and assembly work. We would also like to acknowledge the mechanical design and assembly contributions of Shing Chang, Joy Stafford-Evans, and Chris Rynne and the management support of Dave Green, Gary Galica, and Jim Glynn at PSI.

## References

1. AFRL TacSat-2/Roadrunner Fact Sheet: <http://www.vs.afrl.af.mil/FactSheets/RoadRunner.pdf>
2. See, for instance: <http://www.afcea.org/signal/articles/anmviewer.asp?a=628&z=38>
3. Xilinx Inc., Aerospace and Defense Product Information:  
[http://www.xilinx.com/products/silicon\\_solutions/market\\_specific\\_devices/aero\\_def/index.htm](http://www.xilinx.com/products/silicon_solutions/market_specific_devices/aero_def/index.htm)
4. Xilinx Inc., MicroBlaze Product Information:  
[http://www.xilinx.com/xlnx/xebiz/designResources/ip\\_product\\_details.jsp?key=micro\\_blaze](http://www.xilinx.com/xlnx/xebiz/designResources/ip_product_details.jsp?key=micro_blaze)
5. International Standard ISO/IEC 10918-1: 1994(E), *Information Technology- Digital Compression and coding of continuous-tone still images: Requirements and Guidelines*.
6. I.S. Reed and X. Yu, "Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution," *IEEE Trans on Acoustics, Speech, and Sig. Proc.* 38(10), 1760-1770 (1990).
7. C.M. Stellman, G.G. Hazel, J.V. Michalowicz, A. Stocker, and W. Schaaf, "Real-time hyperspectral detection and cueing," *Opt. Eng.* 39(7), 1928-1935 (2000).