


The Modular Infrared Molecules and Ices Sensor for ESA's Comet Interceptor mission



Liam McSherry,

Principal Engineer, Department of Physics, University of Oxford

Millions of kilometers from Earth, FPGAs are key to helping us learn more about the origins of our Solar System. This is the long journey we're taking to get there.

In 2018, the European Space Agency (ESA) asked the scientific community for proposals for a new 'Fast class' of missions: faster, lower cost, and allowing more experimentation than flagship programs. The selected mission would be a payload of opportunity sharing a launch with the medium class ARIEL exoplanet telescope to the Earth-Sun L2 Lagrange point, around 1.5m kilometers from the Earth.

Comet Interceptor, or Comet-I, was selected for further study in 2019. Travelling at 70km/s past their target, three spacecraft flying in tandem will take the first in-situ observations of a long period comet by imaging and sampling its nucleus and coma (tail). The intercept, which may last as little as seven minutes, will come after up to three years parked at L2 waiting for ground-based surveys to identify a target. By sitting in wait, Comet-I will be able to intercept its target comet before the comet can transit the inner Solar System and before heating from the Sun can reshape its surface, alter its chemistry, and change its temperature. In this 'pristine' state the comet will appear as it did when it was ejected from the edge of the Solar System, offering a time capsule of the material building blocks that became our planets and insight into the processes that form planetesimals. ESA formally adopted Comet-I in 2022 with a planned launch in 2029.

Developing the Modular Infrared Molecules and Ices Sensor

One of the instruments on Comet-I's main spacecraft is the Modular Infrared Molecules and Ices Sensor (MIRMIS). MIRMIS is an infrared imager sensitive over 0.6–25 μm , allowing it to map both the 'thermal IR' light emitted by the comet due to its temperature (8–15 μm) and the spectral lines from minerals (1–2 μm), water (3 μm), methane (3.3 μm), CO₂ (4.3 μm), and carbon monoxide (4.7 μm).

As shown in Figure 1, MIRMIS consists of two independently steered telescopes illuminating three sensors: a Thermal Infrared Imager (TIRI) from the University of Oxford and two integrated Mid-Infrared (MIR) and Near-Infrared (NIR) modules developed by VTT Finland. These are housed in an aluminum chassis with a shared command and data handling unit (CDHU), both built at Oxford. TIRI is a filter radiometer which uses optical filters to divide its field of view into different spectral bands (Figure 2) and will steer its telescope to image the comet through each of the filters, building a full 'cube' of data once it has visited all filters. The sensor behind these filters is a microbolometer array – a grid of 640 \times 480 thermistors coated with a light-absorbing material where the resistance of each pixel, used to set a current and measured with a transimpedance amplifier, is related to the intensity of its illumination.

Data from TIRI's detector can be applied in many ways. Most simply, like a conventional camera, thermal images of the comet will show its size, shape, and surface temperatures. Used with spectral knowledge from TIRI's filters, scientists can fit the light intensity to the well-known spectral lines of various chemicals to estimate their abundance and distribution over the comet, giving an idea of the chemistry of the early Solar System and of comets before they encounter the Sun.




Figure 1: A functional model of the MIRMIS instrument for ESA's Comet Interceptor mission.




Figure 2: Simplified view of MIRMIS-TIRI's optical architecture.

Combined with information about TIRI's angle relative to the comet, it will enable scientists to analyze changes in how light is reflected to make estimates about the surface topology of the comet and the size and presence of dust and debris, offering added insights into the comet's formation that the Solar wind would have otherwise blown away.

Where FPGAs fit into the picture

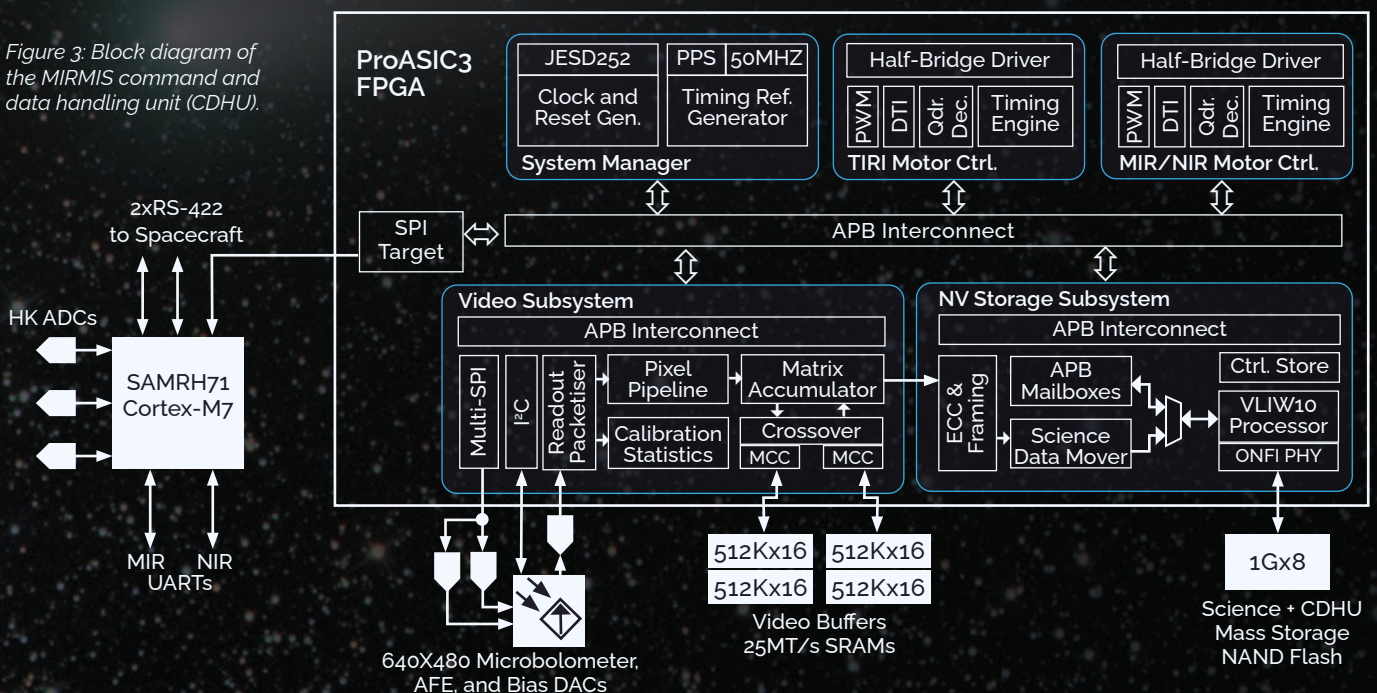
Within MIRMIS, TIRI and the CDHU are closely integrated (Figure 3). The CDHU has a radiation-hardened Arm Cortex-M7 microcontroller that communicates with the spacecraft, monitors low speed housekeeping sensors, performs trajectory calculations for motor pointing, and manages high level scheduling for both the TIRI and MIR/NIR segments. The CDHU also has a radiation-tolerant ProASIC3 FPGA that deals with hard real-time or high data rate tasks such as pointing motor control for the telescopes, capturing and processing images from TIRI's detector, and operating the 8Gbit NAND flash array used to store both the configuration data for the CDHU microcontroller and the science data captured from TIRI's detector. The MIR/NIR segment operates largely autonomously, only receiving high level scheduling from the CDHU.

The architecture of TIRI and the CDHU is a natural evolution from the University of Oxford's Lunar Thermal Mapper (LTM). As imaging sensors grow both in resolution and frame rate, and while processors for deep space stay relatively constrained compared to those available in the near-Earth environment, it makes sense to move more functionality into an FPGA and to simplify the flight software.

In MIRMIS, evolution from LTM has resulted in the integration of multiple programmable elements so that almost the entire image capturing process is offloaded from the CDHU microcontroller. The TIRI and MIR/NIR motor controllers include 'timing engines': small processors that, based on the trajectory computed by the CDHU microcontroller, tweak motor speeds and positions and trigger image captures with 1µs precise timing synchronized to a spacecraft-wide pulse per second (PPS) signal. Requested motor positions are achieved with closed loop stepper motor drivers that monitor either a quadrature encoder on each motor's shaft or, as a fallback, the number of steps the driver has generated. Once triggered by a timing engine, the CDHU FPGAs video subsystem in turn triggers TIRI's detector to begin reading out a frame.

In the leadup to comet encounter, captured frames will be processed inside the CDHU FPGA to measure statistics about them, such as the average pixel value and a histogram over the frame. The CDHU microcontroller uses these statistics to tweak four gain and integration parameters for TIRI's detector to bring the pixels to mid-scale. By imaging both a space view and an onboard black body for calibration, TIRI retains traceable sub-Kelvin temperature measurements despite the adjustments made to the detector. During encounter, the video subsystem is autonomous – its readout packetizer grabs pixels and temperature data from the detector and formats them into a packet with other metadata such as motor speed and position.

Figure 3: Block diagram of the MIRMIS command and data handling unit (CDHU).



These packets flow into the pixel pipeline where simple conditioning can optionally be applied, including dithering to spread quantization error and bit shifting to reduce the magnitude of very bright pixels. Once conditioned, the packets enter a matrix accumulator that can sum consecutive packets to improve signal to noise ratio, and apply rounding and saturation to reduce error. The matrix accumulator is double buffered, allowing TIRI to capture and accumulate new packets while the previous packet is drained into non-volatile storage.


The CDHU FPGA's non-volatile storage (NVS) subsystem is similarly autonomous. A packet that exits the video subsystem first has an error correcting code (ECC) applied and is then framed for storage, serialized from its internal wide format into the octets accepted by the NAND flash, and passed to an Open NAND Flash Interface (ONFI) controller and PHY. The controller is a fixed function state machine that tracks occupied regions of the flash, transfers serialized packets into the array, and initiates PROGRAM PAGE commands, while the PHY is a deterministically timed I/O processor that generates the waveforms needed to operate an ONFI device based on the controller's commands. Once a packet is stored, the video and NVS subsystems return to idle until next triggered. Throughout this process, the only intervention required from the CDHU microcontroller is to load new commands into a timing engine based on trajectory.

This data transfer within the CDHU FPGA uses a separate AXI-Stream path, and each of the components has dataflow-based operation. The result is a high speed data path that can operate at the full speed of the NAND array but, within which, brief stalls are handled gracefully by the bubbling up of backpressure and the 'distributed FIFO' formed by many components each with their own small internal storage. Following encounter to retrieve science data, and during instrument boot to retrieve configuration data, the CDHU microcontroller bypasses this AXI-Stream path and uses a set of mailboxes exposed through an APB interconnect to communicate with the I/O processor. Similarly, configuration registers for all functional units are attached to the APB interconnect that the CDHU microcontroller accesses through an SPI to APB converter.

The importance of validation and verification

The CDHU FPGA's image capturing data path in particular, but also more broadly the other parts of the CDHU, are a complex and interconnected system. Verification is the largest slice of effort in their development and is driven mainly by the European Cooperation for Space Standardization (ECSS) standards, which mandate extensive and traceable requirements and documentation. Validation, however, is just as important: in a harsh environment such as deep space, behavioral or gate level simulations cannot give the same confidence they do on the ground. A system must be able to detect and recover from, for example, flipped bits or locked up devices caused by charged particles striking the electronics, analogue components degrading under radiation, and the failure of mechanical assemblies from shock or harsh vibrations. A functional design is necessary but not sufficient for success – it must also be the correct design for space.

Within the CDHU, practically all functional units are influenced by this thinking. A watchdog on the the APB interconnect monitors for bus faults and the SPI target can report these faults whether or not the bus is functional. If the SPI target fails, a separate JESD252 block can reset the core logic in the absence of higher level communication. The PPS reference, as well as having redundant inputs from the spacecraft, can automatically fall back to a counter driven by the FPGA's logic clock.



The motors have redundant zero position encoders and the controllers can operate closed and open loop so that minimally imprecise pointing is possible even if an encoder should fail. The I²C controller does not implement clock stretching and so cannot be stalled by a misbehaving target.

The I/O processor is an example worth calling out: it sounds complex, but the fully programmable timings mean that a degrading NAND array can be compensated for – although, practically, it also means that logic depth and routing density can be reduced on what is already a slow fabric by avoiding encoding timings in a large state machine. The processor itself is small (≈ 200 LUT3+FF) and takes a 512 instruction microprogram, stored in a RAM with ECC and a memory scrubber to correct bit flips, with much of the complexity delegated into the compiler that only runs on the ground.

In other cases, the need to change a design is not directly due to the space environment. Most synthesis tools can automatically implement local triple modular redundancy (LTMR), where registers are triplicated, compared, and corrected. This at first appears to be a solution that needs little thought, but the resulting increases in resource use, routing congestion, and logic depth – taken with the performance of the 20-year-old ProASIC3 – mean that decisions inconsequential on a modern FPGA can seriously degrade timing performance. As an example, the lack of shift register LUT (SRL) equivalents on the ProASIC3 means that even small shift registers have a very large area, and with LTMR enabled it can become challenging to run a single cycle 5 bit adder at 100MHz.

Verification for MIRMIS' CDHU involves an extensive requirements derived test suite. Within the FPGA, functional units are divided into their smallest parts and, where tractable, each undergoes a bounded model check against a formal proof.

The same proofs are reused at the integration level, only replaced by conventional testbenches for very long simulations or for tests involving complicated calculations. Development follows an iterative, spiral approach with regular releases to MIRMIS' software team, and the FPGA and flight software undergo automated hardware in the loop testing in their flight configurations. The functional coverage provided by formal verification and hardware in the loop testing has allowed an extremely fast start to finish development time of around 18 months with very little rework.

Next steps

These instruments continue to be refined and improved. In March 2025, the University of Oxford demonstrated its Broad Horizons prototype, an evolution of TIRI with a widened 25×12° field of view and a 1280×1024 microbolometer array detector that will be the basis for a Mars observing instrument submitted to ESA's Lightship-1 call. LTM and MIRMIS derived instruments also form parts of proposals into ESA's M8, F3 and mini-F calls.

This work was supported by the UK Space Agency's Space Science Programme and, in part, by a grant under the Centre for Earth Observation Instrumentation's Pathfinder Small Projects programme. Formal verification tools were provided at no cost by YosysHQ GmbH.

