ADIUVO
ENGINEERING AND TRAINING, LTD.

AVNET®

/ SPACE APPLICATION CASE STUDIES

TABLE OF CONTENTS

INTRODUCTION

Space plays an increasingly important part in our lives. Not only does it provide new discoveries and science though game changing missions like James Webb, Plato and Lunar Return, space technologies also enable modern communications, modern banking, navigation, and the ability to monitor the health and security of our planet.

Satellite communications have enabled connectivity not only while we travel on ships and aircraft, but have also opened up communications in previously geographically or politically inaccessible areas. Satellite navigation provided by global systems such as GPS and Galileo and regional systems such as those covering Japan and India enable accurate navigation across the globe and also provide a precision time reference. This time reference is the backbone of most modern Internet-based transactions, as it enables timestamping of transactions, which allows the ordering of such transactions. When it comes to health and security monitoring of our planet, governments and international organizations operate Earth Observation Satellites (EOS), which monitor everything from weather, wild fire, vegetation, atmospheric pollution, and ocean health.

Over the last 10 years, space has seen rapid growth in companies providing a range of capabilities from launch to satellite design and deployment, along with managing the services and data provided from space. Downstream data is seen as increasingly valuable and satellite operators often go to great lengths to protect it. The value of this data can be financial – for example, enabling accurate tracking of assets in transit across the world, to securing potential Nobel prize winning discoveries from exploration missions.

These satellites use a variety of sensors and communication methods, from image sensors for earth observation and scientific exploration to RF for communications satellites, Synthetic Aperture Radar Imaging and laser communications between satellites in different orbits.

Satellites are complex systems of systems whose functionality and mission is enabled by the combination of several different systems each performing a vital role. Satellite architectures are generally split between the platform and payload. The payloads are the systems which implement the functionality of the mission, such as EOS or communications. The platform provides all of the services necessary to support the payload (for example, power and battery management, station keeping, telecommand and telemetry, thermal control and housekeeping, etc.). Each of these systems come with their own challenges, from interfacing to reliability, security, and the traditional size weight and power (SWaP) challenges.

FPGAs play a role at the heart of many of these space applications thanks to several of their architectural features:

- Performance – FPGAs enable parallel structures to be implemented, supporting high performance processing chains necessary for applications such as image and signal processing and communications. Programmable logic also enables developers to be able to implement radiation mitigation features such as majority voting, error detection and correction codes on memories, etc.

- Interfacing – FPGA IO structures are very flexible. They are able to implement many IO standards directly (such as LVDS, LVCMOS, Gigabit Transceivers, etc.). When combined with an external PHY (e.g. Mil Std 1553), they are able to implement any IO standard which may be required. This is critical for space applications, as some missions demand high performance IO to support high throughput interfaces. While other applications require a designer interface

AMD KINTEX™ ULTRASCALE™ FPGA-BASED IMAGER FOR EARTH OBSERVATION

One of the main uses of satellites is earth observation. Earth observation is conducted for a number of reasons, from observing planetary health to monitoring shipping and port activity, farming and other land use, intelligence, surveillance and reconnaissance conducted by nation states. Earth observation can take place using several different sensor modalities depending upon the application, from synthetic aperture radar to visual and multispectral imaging.

Multispectral imaging uses not only visual wavelengths, but also elements of the near infra red band. Typically coatings are placed on the CMOS imaging sensor itself to filter out all but the wavelengths of interest. This enables space-based monitoring of chlorophyll and aerosols for plant health.

CMOS image sensors, which are used in industrial and commercial applications terrestrially, tend to use 2D image sensors to create a 2-dimensional image. However, space-based earth observation image sensors use what is a called a line scan sensor (often a Time Domain Integration), which is a long sensor array that constructs a two dimensional image as the satellite passes over the earth in its orbit. This means that line scan sensors often produce a great volume of data when imaging. This data has to be received quickly and processed and stored for downloading for later analysis.

It is typical that an EOS satellite will not be continually imaging, but will be turned on to image targets of interest as needed. When a TDI imager is operational and capturing images, the quantity of image data generated can be significant. The data volume generated depends on the sensor configuration, orbit and configuration. To provide some idea the volume of data, a TID CMOS sensor travelling at 7 Km per second would generate 20 Gbps of data across its multispectral channels. This means a 1000 KM imaging swath will produce 2840 Gb (355GB) of data. The data volume depends on the sensor type used and its exact configuration. This data has to be received, recovered and packaged such that it can be forwarded to a Data Handling Unit (DHU) for storage in system mass memory.

The CMOS image sensor (CIS) will provide the output for each of its multispectral channels using Current Mode Logic serial links.

Creating a front end image processing system that is capable of interfacing with the CMOS sensor presents the system designer with a lot of challenges. Not only does this data need to be received, processed and forwarded to the DHU, the solution must also be low noise to ensure there is no ability for the noise from the power system to impact the performance of the CMOS image sensor. The Front End Electronics must also be able to provide the CIS with low jitter clocks along with configuring sensor parameters and operation over a common link such as SPI or I2C. The Front End Electronics (FEE) electronics may also need to be able to adjust sensor voltages and biases during the mission to ensure radiation and aging effects are mitigated and image quality does not decrease.

Along with the formatting of the received image data, it also necessary for the FEE system to tag the image information with metadata. This metadata might include elements such as telemetry, including temperatures and voltages, but also spacecraft information such as time and position. Including this metadata in the information at the front of the processing chain in the FEE ensures the necessary metadata flows through the processing chain, included within the image. This means that the FEE needs to have interfaces to spacecraft systems. Which interface is used will vary from system to

system, however, it may be interfaces such as CAN, SpaceWire, Mil Std 1553, etc.
The context diagram for the FEE can be seen below.

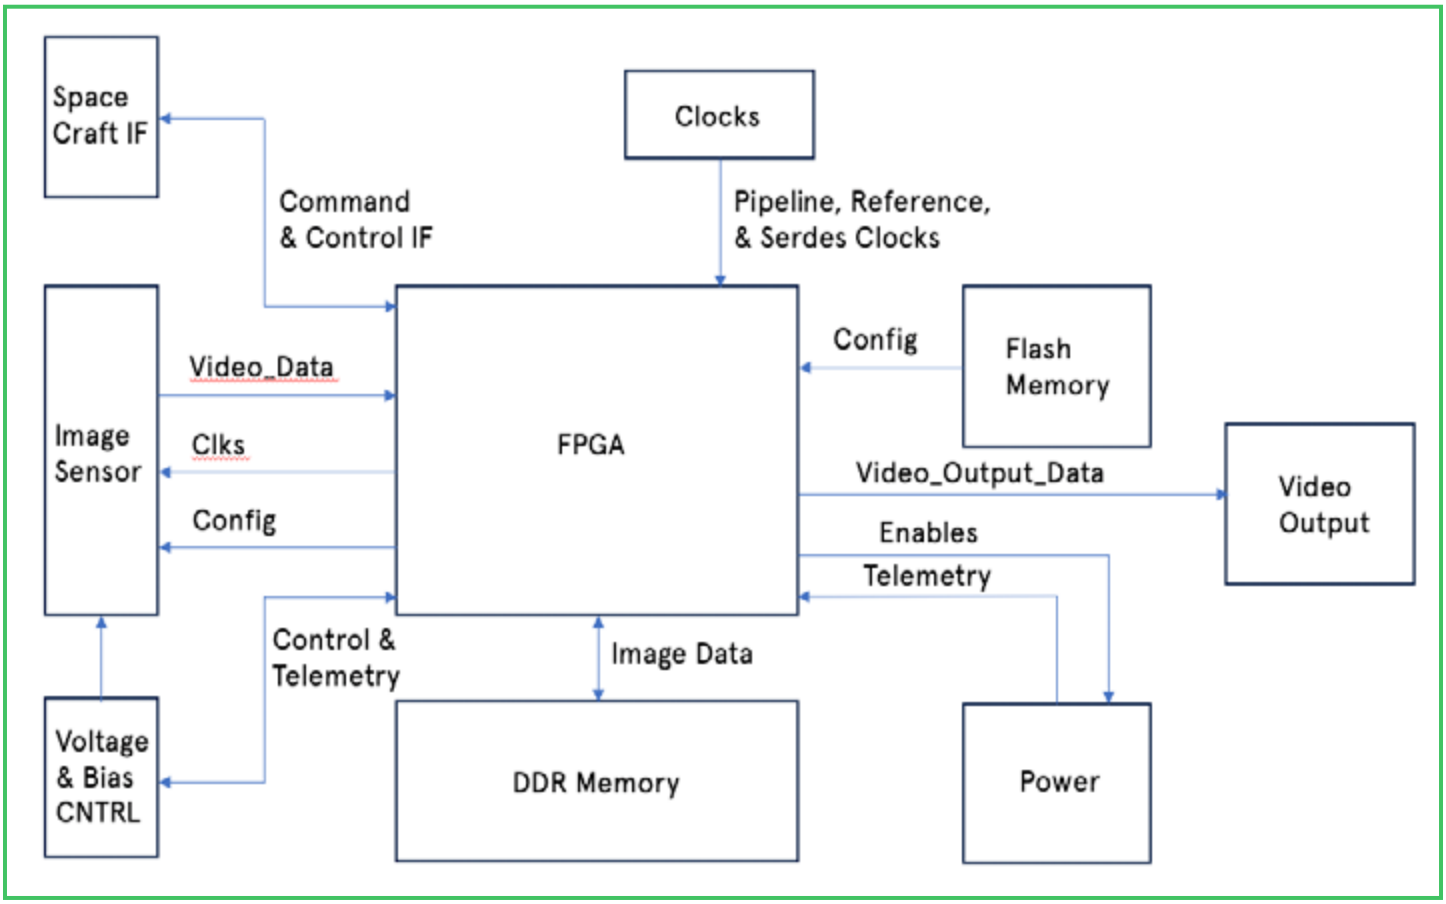

To address these challenges presented by an FEE, an FPGA design using the AMD Kintex XQRKU060 FPGA is a good approach. It provides the flexible interfacing needed to interface with CMOS image sensor leveraging the multi gigabit transceivers (MGT). These MGTs can also be leveraged to connect the FEE to downstream processing units such as the Data Handling Unit (DHU). Due to the large number of GTH available on a Kinex UltraScale, a prime / redundant FEE to DHU interface maybe considered depending upon the mission reliability requirements.

The flexibility of the IO cells will also enable the interfacing with spacecraft computers using interfaces such as SpaceWire, CAN, and 1553, etc. This enables the sequencing of the imaging operations to be commanded and controlled by the spacecraft. It also enables the reception of some of the necessary metadata.

The flexible clocking architecture of the AMD Kintex UltraScale device with its provided Mixed Mode Clock Managers (MMCM) enables the design in the AMD Kintex to be able to generate all of the necessary reference clocks required by the CIS Imager. The flexibility of these clocking resources enables the reference clocks to be adjusted in real time to support different CIS Imager configurations.

Depending on the architecture of the payload, the FEE may also be required to provide thermal control of the CMOS image sensor. The FPGA logic can be used to implement the necessary monitoring and drive structures, which the Triple

Modular Redundant (TMR) AMD MicroBlaze™ soft processor can be used to perform the necessary thermal control (Proportional Integral Differential, or PID) algorithms.

The parallel structure of the programmable logic combined with DSP elements and Block RAM is ideal for the development of high performance image processing algorithms. This image processing chain can start with the reception of the data and perform any necessary image processing algorithms (e.g. Co Registration of Bands, Super Resolution). This may also require the use of external high performance memory such as DDR4, which is supported by the AMD Kintex XQRKU060 FPGA. Internally, the architecture of the FPGA can leverage standard high performance interfaces such as AXI and AXI streaming to interconnect the image processing pipeline. Developers are also able to leverage the large IP library provided by the AMD Vivado design suite.

This image processing pipeline could also be implemented using model based approaches to accelerate the development timescales. For example, the MathWorks MATLAB Simulink® and HDL Coder can be used to quickly develop image processing elements which can be deployed in the AMD Kintex UltraScale FPGA very quickly and easily.

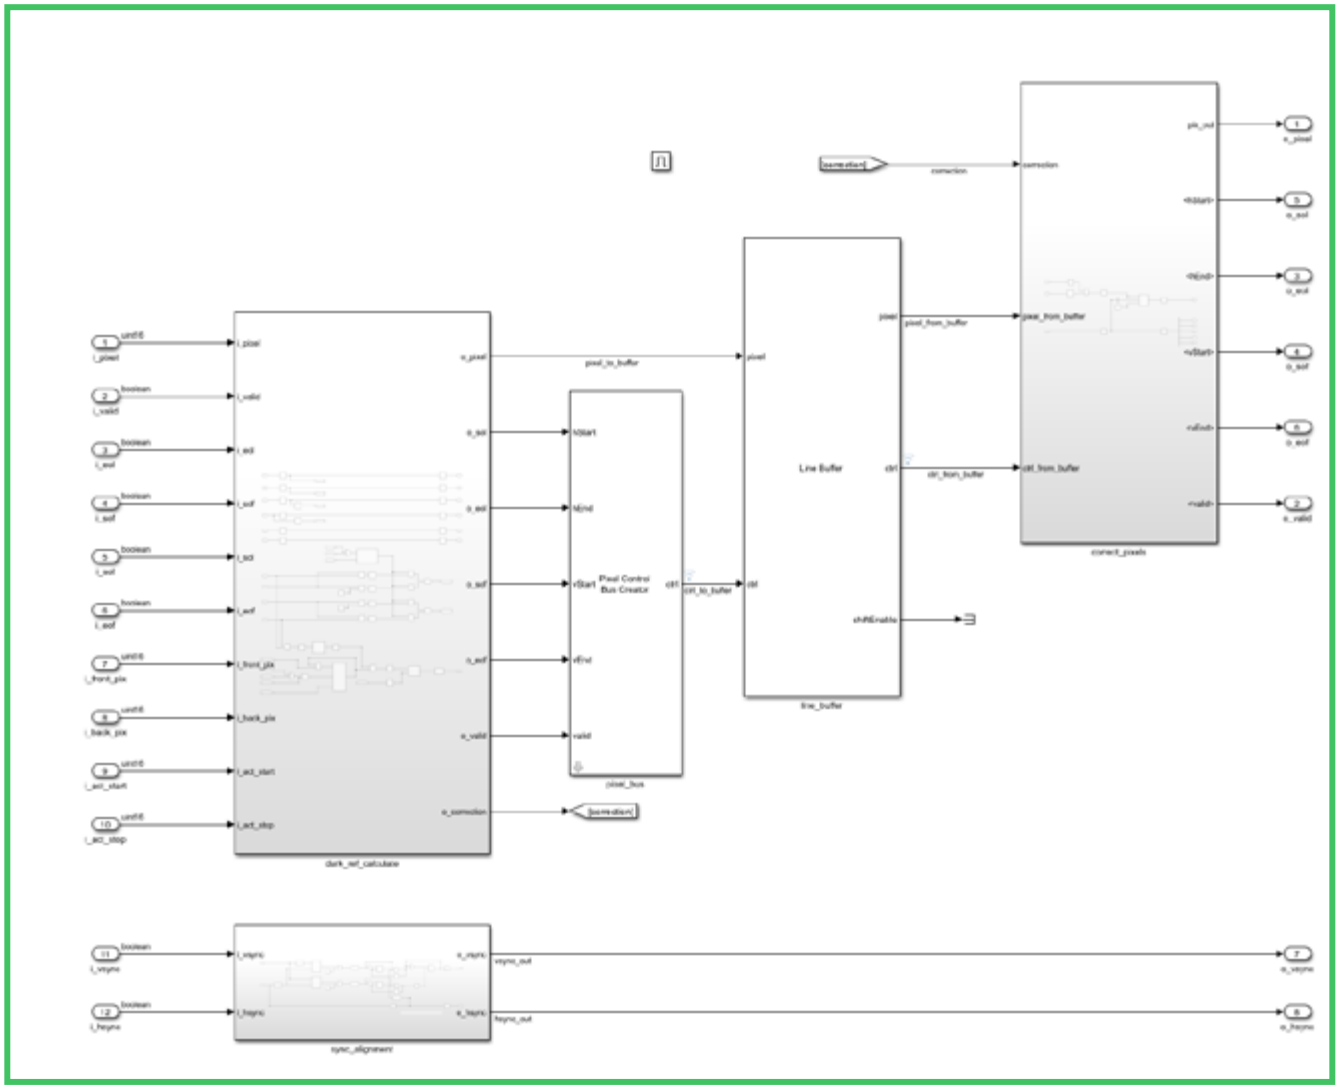

MathWorks MATLAB Simulink implementation of dark pixel subtraction

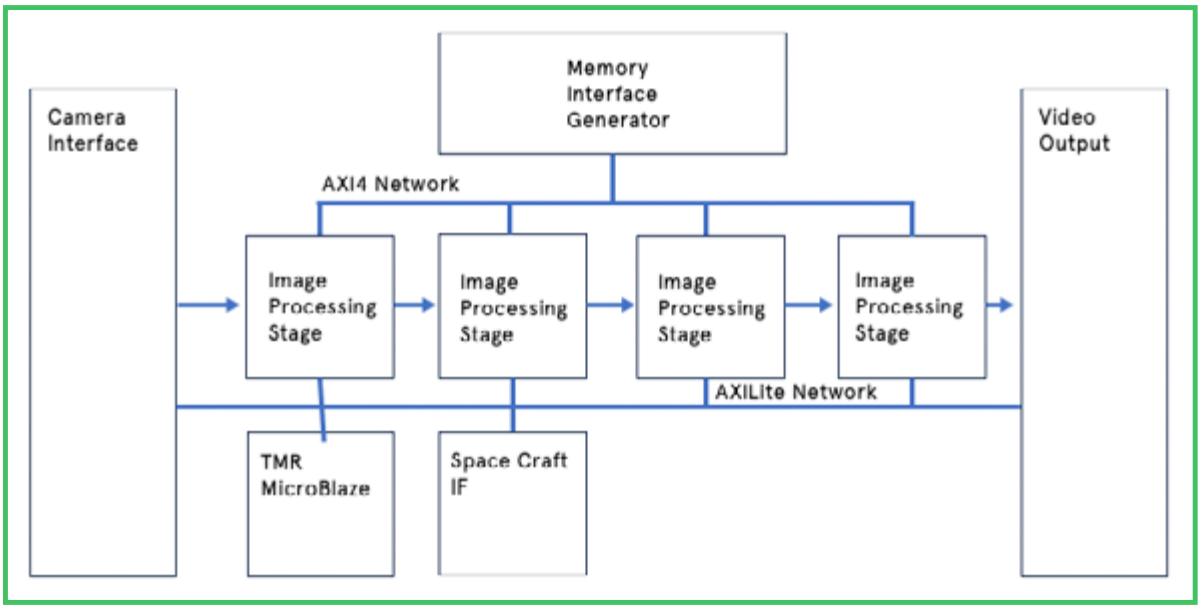

The use of a AMD Kintex XQRKU060 FPGA to implement an FEE unit for working with high performance TDI CMOS Image Sensors provides developers with a flexible solution that addresses the challenges presented when working with such complex image sensors.

AMD KINTEX ULTRASCALE FPGA-BASED EDGE AI PROGNOSTICS

Equipment failures on satellites, regardless of cause, can create significant impacts on the operation of the mission. This could lead to a distant exploration satellite going into a safe mode and stop performing its scientific mission or a telecommunications satellite to drop its communication channels.

Satellites are designed with considerable redundancy and the removal of single points of failure. This ensures that if random or systematic failures occur, the overall mission and desired operational lifetime can still be achieved.

Satellite systems have complex telemetry required in order to understand the health of the satellite and if random or systematic failures have occurred. Depending on the system, the telemetry provided by the systems will vary from simple secondary voltages, secondary currents and board temperatures to more complex elements such as solar array currents, battery temperatures, battery voltages, power switch states, overcurrent and brown out information, operating modes, memory errors / corrections, Attitude Determination And Control System (ADCS) information, etc.

Typically this telemetry information is transmitted to the ground station, enabling the operations team on the ground to analyze the data, make appropriate decisions, and reconfigure the satellite systems. Depending on the location of the satellite, this information could be relatively recent or delayed, which also makes issuing of corrective telecommands slower.

In some cases, satellite systems need to be able to take corrective action without the appropriate telecommand to ensure critical functionality is not damaged should a failure occur. One example of this would be the front end electronics temperature going beyond allowable limits. In this case, it would be expected for the FEE thermal control system to monitor and detect the base plate temperature excursion and take an appropriate action to protect the FEE.

While these systems protect the satellite, it is much better to be able to perform prognostics on the telemetry data. Being able to identify systems which might be beginning to experience issues earlier can enable the operations team or the satellite to take mitigating actions before the failure occurs.

FPGAs are also ideal for the implementation of machine learning (ML) inference thanks to the parallel structure of programmable logic.

ML is ideal for prognostic applications, as it enables the data to be analyzed, trends identified, and issues corrected before they become problems. Another benefit of the FPGA is their flexible IO structures, which enable the interfacing with a wide range of sensors, which are used to gather the data.

Unlike many ML applications, high throughput is not required, as telemetry information is normally gathered at a very slow frequency (e.g. 1Hz).

One method of implementing prognostics within a FPGA is to implement an ML inference using TinyML algorithms that are running on a softcore processor such as a MicroBlaze. Many applications that use FPGAs often use a softcore processor implemented within the device to manage the sequential communications, gather telemetry and run the payload configuration.

For example, an image processing front end electronics system which interfaces with and processes data from a CMOS image sensor will provide telemetry on the focal plane temperature, supply voltages and bias voltages along with the FEE voltages and currents. This would be a typical application use case for an AMD XQR Kintex UltraScale FPGA.

A TinyML implementation could be deployed on a MicroBlaze processor. TinyML application would be able to analyze the telemetry from the image sensor supplies to identify changes in the sensor voltages and currents. Changes to these voltages and currents would indicate a parametric shift in the image sensor due to radiation damage. Running an ML model would enable prognostics to identify the aging of the sensor based off these models and, if necessary, make the necessary adjustments to the supply and bias voltages to correct for performance.

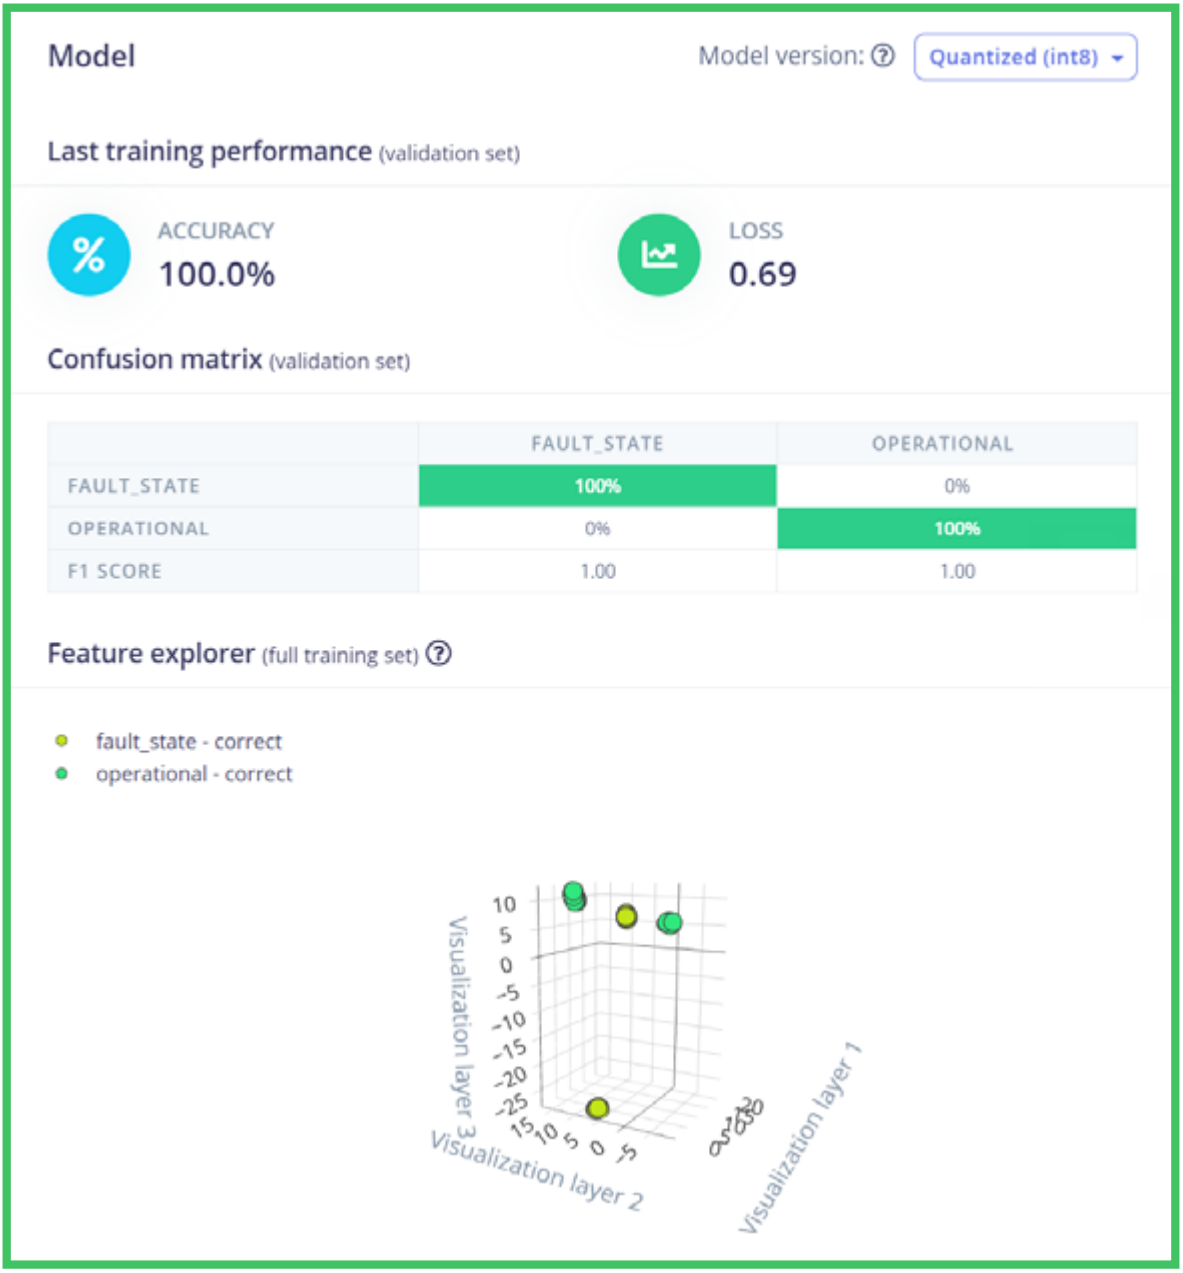

TinyML model identifying fault state depending on current draw

AMD KINTEX ULTRASCALE FPGA IMAGE PROCESSING FOR SCIENTIFIC EXPLORATION

In a similar manner to earth observation satellites, several satellites such as the Hubble Space Telescope are used to observe the universe around us, gathering data for new scientific discoveries. Unlike earth observation satellites, which use line scan imagers, scientific missions often use 2D CMOS Image Sensor (CIS) or Charge Coupled Devices (CCD) to capture the images. Depending on the mission, there may be a single image sensor or an image sensor array that allows the application to capture different wavelengths and aspects enabling discoveries.

When working with 2D image sensors, a range of processing is required on the sensor data to ensure an optimal noise free image. The first challenge is the driving of the sensor. If a CCD sensor is used, precise timing must be implemented to enable the charge to be transferred without loss in the sensor. However, for CMOS sensors, complex configuration sequences may be required to configure the imager for the correct imaging mode.

For both CCD and CIS sensors, the pixel data is made available using either high speed differential LVDS or Gigabit Serial Links either directly from the device (CIS) or via an ADC first (CCD).

However, the processing pipeline for these imagers can be extensive, as it includes features such as reference frame subtraction, dark reference correction, pixel correction, cropping and binning.

These image processing pipelines require high throughput to perform the required processing while maintaining the frame rate required.

Several of these operations such as reference frame subtraction and dark reference correction require reference frames to be either stored in external memory following initial commissioning and calibration of the sensor or gathered and then stored within external memory during the operation.

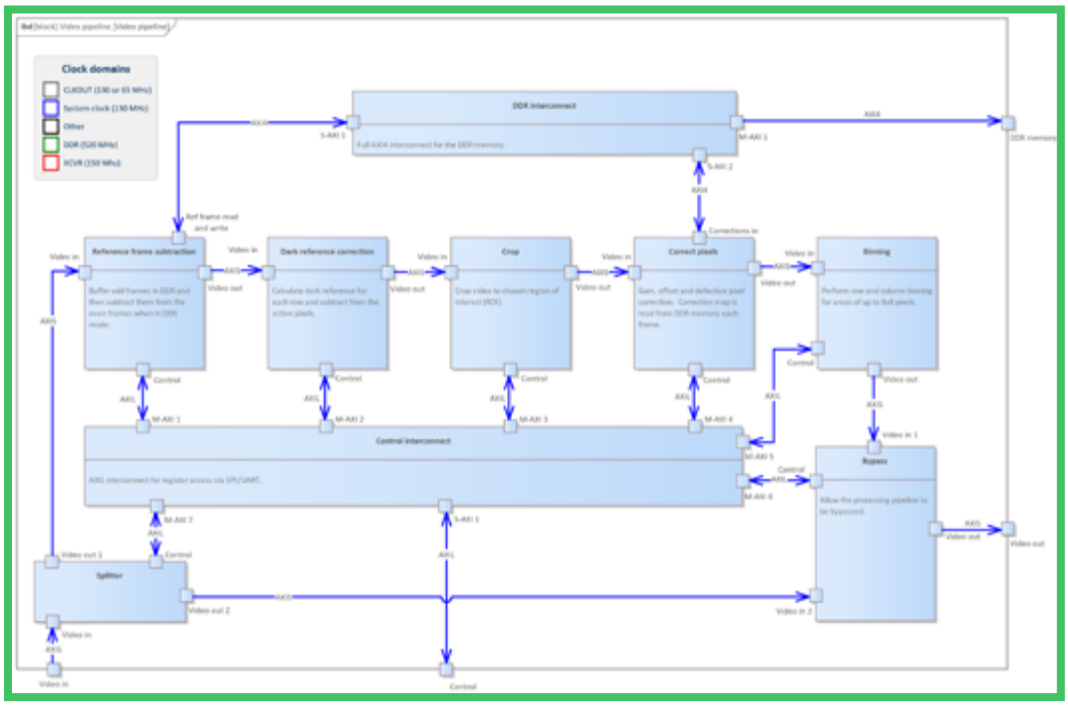

The output from this image processing pipeline will then need to be correctly formatted and packetized for communication to the downstream equipment. This formatting could be SpaceWire, SpaceFibre or a networked protocol such as GigEVision. The standard used will vary from mission to mission.

For scientific imagers, AMD XQR Kintex UltraScale FPGAs are an ideal choice as they enable the implementation of the high speed image processing pipeline. The AMD Kintex UltraScale is able to support DDR4 memories which enables the image processing pipeline to store and retrieve images as necessary as part of the image processing pipeline, while the 726K LUTs, 2760 DSP elements and 38Mb of embedded memory can be used to implement image processing features.

The flexible nature of the the AMD Kintex UltraScale FPGA IO enables the device to be able to interface with sensors outputs which use LVDS or similar standards. Gigabit Serial Links can also be used on the AMD Kintex UltraScale FPGA to support image sensors which have a higher data rate.

Control of the imaging device, its voltages and biases, configuration of the image processing pipeline, house keeping and telemetry can be provided by the implementation of the triple modular redundant MicroBlaze processor. These TMR MicroBlaze processors can also be responsible for command and control communication with different satellite modules using interfaces such as SpaceWire.

TMR MicroBlaze instantiates three MicroBlaze softcore processors and runs them in tight lockstep. Each operation on its Data and Instruction buses are voted on and corrected to ensure high reliability operation. Design features such as isolation flow can be used within the design implementation stages to physically separate the three instances within the device.

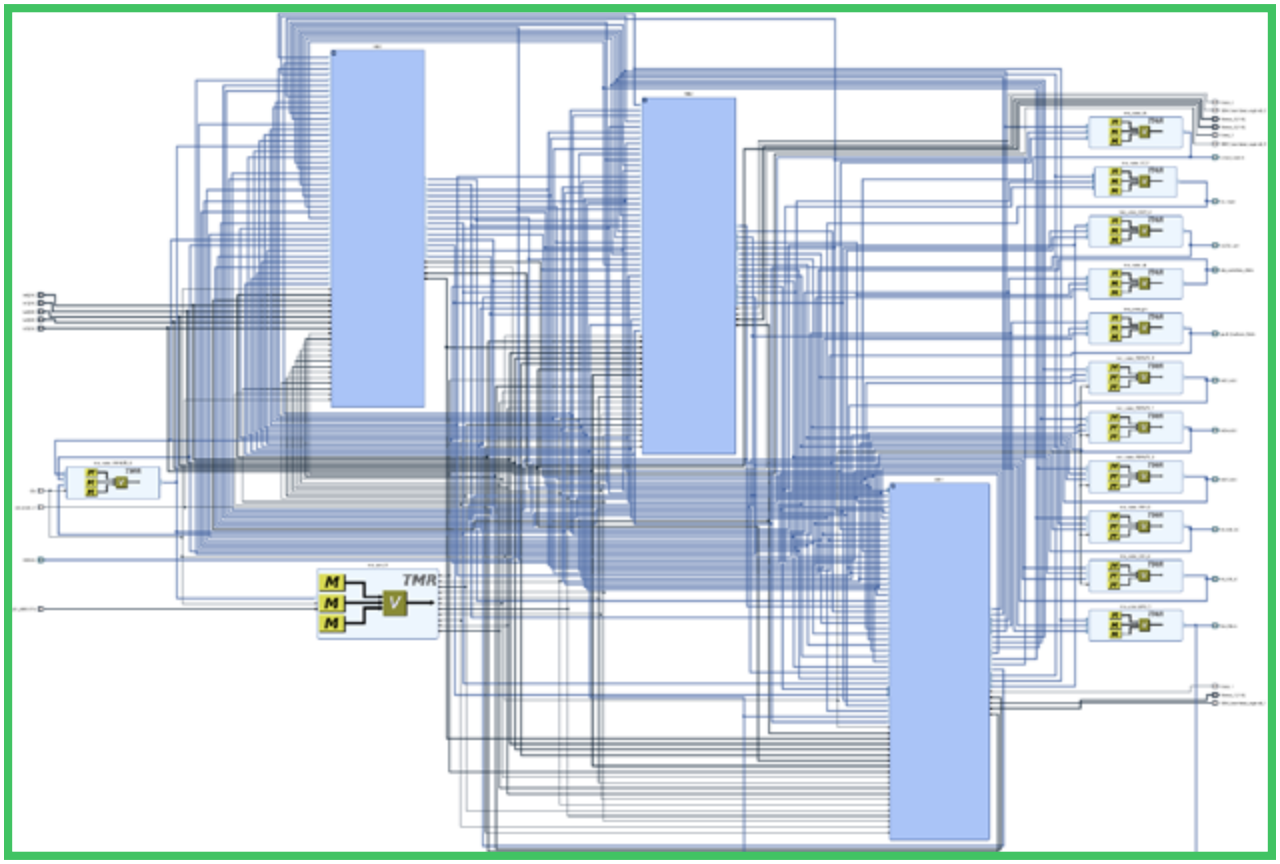

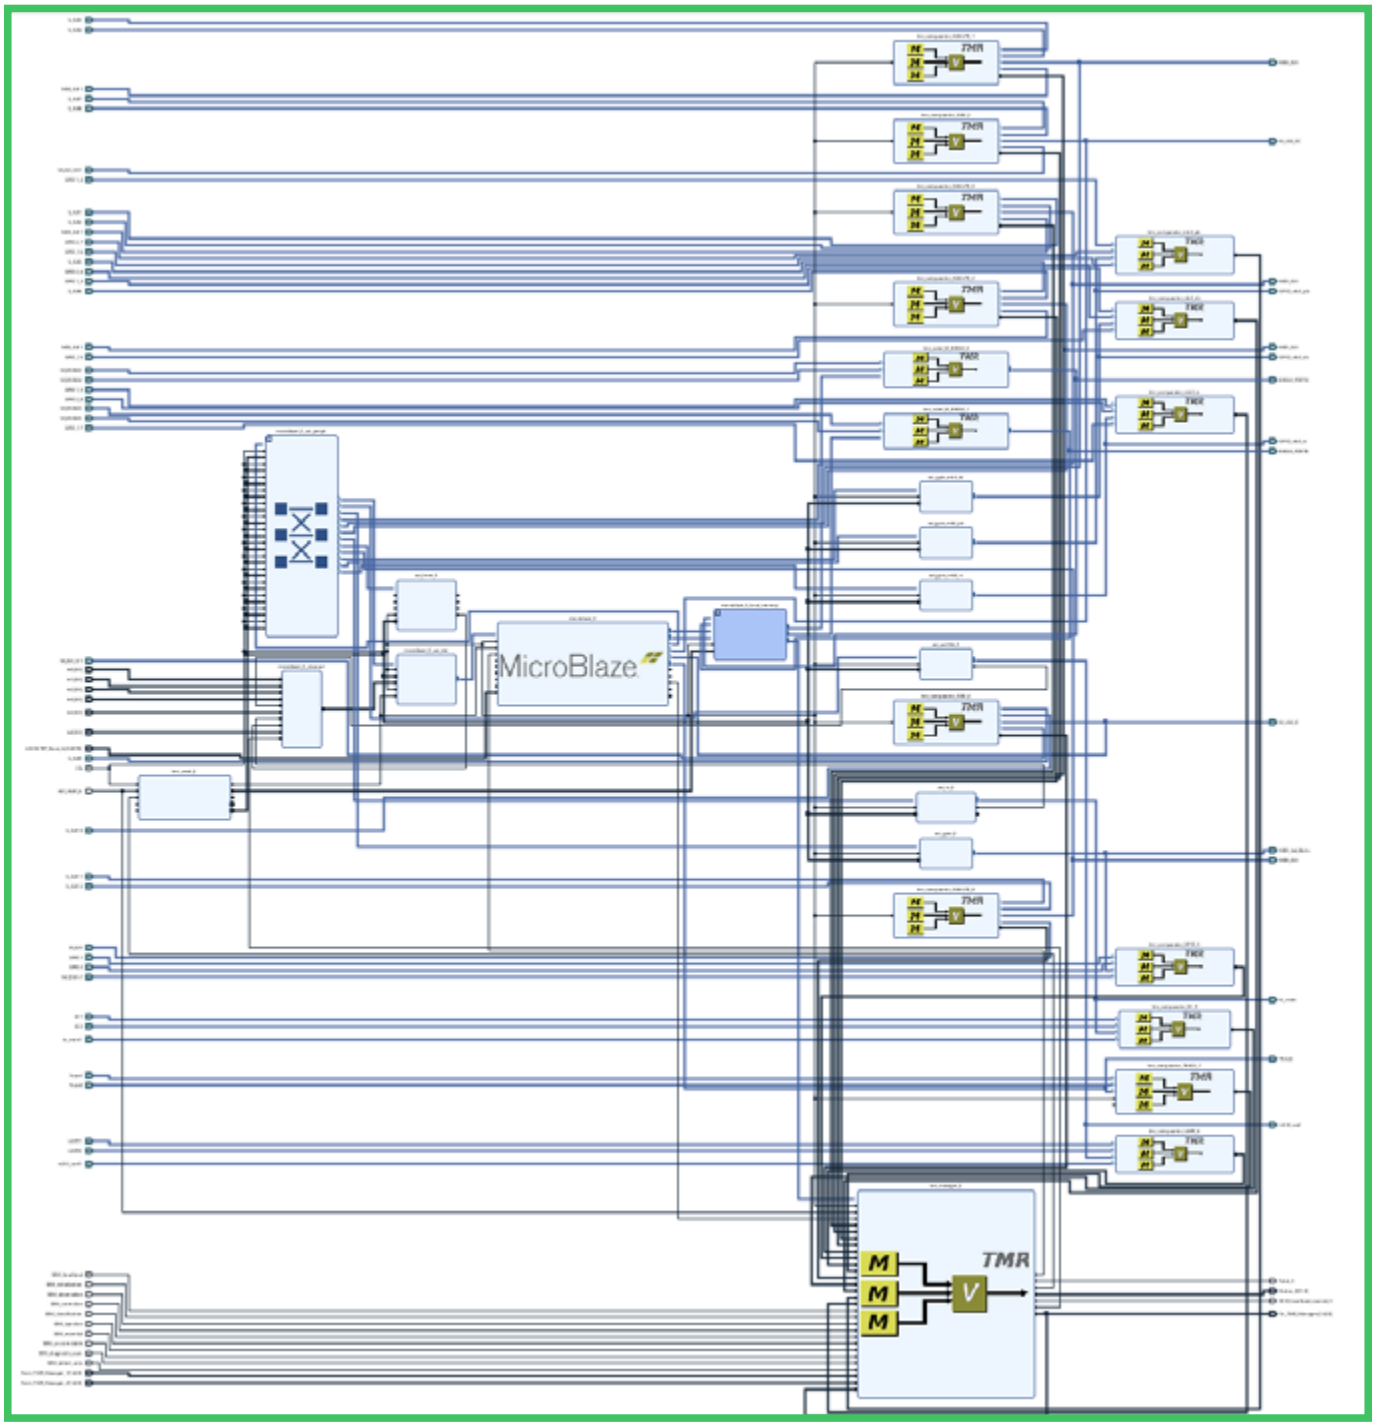

Because of the large logic capacity of the device, a neural network can be deployed on the the AMD XQR Kintex device with ease. These neural networks can be used to classify images received from the sensor arrays to identify objects of interest or image quality (e.g. cloud or weather cover). Another use case of machine learning is estimating the radiation damage occurring to the sensor based on changes in quality of the images. This would enable the image processing sensor and its drives and biases to be adjusted to ensure continued performance of the sensor as it ages in operation. Applications such as this would be developed using the Deep Learning Toolbox and Deep Learning HDL Toolbox from MathWorks which provide the ability to implement neural networks.

AMD VERSAL™ ADAPTIVE SOC INTEGRATED IMAGE FRONT END ELECTRONICS, AI AND DPU

Like many applications, satellites often face size, weight, and power, and cost challenges (SWaP-C). This is especially true when several systems are required to implement the overall mission. For example, a typical Earth observation satellite will have several different systems which perform differently, but critical functions to gather an image, and process it such that it can be transmitted to the ground station on demand.

A typical Earth observation satellite image chain would comprise of the following elements:

- Front End Electronics (FEE) – This configures and drives the image sensor, recovers the image and post processes it for transmission onwards.

- Data Handling Unit – The DHU receives the images from the FEE performing additional processing on the image frames received. This includes loss less compression, file formatting and storage in mass memory. The DHU also has interfaces to the satellite computer to retrieve images from mass memory and forward them on for transmission to the ground station. To make the valuable bandwidth used as efficiently as possible, DHU may also examine image frames and classify them on quality for example using machine learning to classify images (e.g. cloud coverage reducing target image quality).

- Thermal Control Unit (TCU) – The TCU will monitor the temperature of the optical telescopes and the FEE electronics to ensure the optical performance is optimal for the imaging performance and is operating within acceptable and safe temperature range.

The FEE, TCU and DHU will each use at least one high performance FPGA. The DHU will also include a high performance processor which performs the communications with the satellite computer and manages the file system in mass memory. The FEE, TCU and DHU will be powered via a complex power chain, which includes isolating dc/dc convertors and then point of load regulators to provide the voltages required by the FPGAs, processors, memories and supporting circuits. Within the satellites themselves, the equipment must be connected together with cabling between the FEE, TCU and DHU, to and from the satellite controller and the power and telemetry systems.

Integration of the FEE, TCU and DHU would address many size, weight and power, and cost challenges which are faced by satellite developers. The integration of the FEE and DHU brings with it significant advantages, including:

- Reduction in cabling necessary for the separate TCU, FEE and DHU. This can result in a significant reduction in costs associated with complex cabling solutions. The reduction in cables required also reduces a significant contributor to the mass budget of the satellite.

- Power Budget Improvements. Power losses through each system isolating dc/dc convertors can be considerable (e.g. 10-30% range depending on convertor efficiency). Each unit then also has to drive its processing core, interfacing and external communication drive. Reduction in bus convertor power losses, coupled with a lower overall power dissipation can provide a significant benefit to the power budget of the mission.

- Reduced EMC / EMI. Integration of the FEE, TCU and DHU enables several processing functions which would previously have been in different components to be integrated within a single device.
- Reliability. With a reduced component count, reduced cabling and a reduced number of Isolating dc/dc power supplies, the overall reliability of the system will be improved.

The AMD XQR Versal adaptive SoC provides designers with the ability to integrate the FEE, TCU and DHU within a single system thanks to its diverse adaptive SoC architecture.

The AMD Versal adaptive SoC architecture provides developers with dual-core A72 processors and real-time dual core R5F processors coupled with vector processors and programmable logic. Connectivity between these elements uses a network-on-chip to ensure the highest performance. In addition to the different processing elements, the AMD XQR Versal adaptive SoC also provides developers with a range of integrated IO peripherals, including standard interfaces such as QSPI, SPI, CAN, UART, GigE, SATA, etc. The AMD XQR Versal adaptive SoC can integrate with a range of volatile and non-volatile memories including DDR4, eMMC, along with high performance memories such as SATA by leveraging the multi gigabit transceivers.

The XQR architecture can therefore be segmented to provide the necessary functionality required to implement FEE, TCU and DHU functionality.

Of all the functionality required, the FEE is the most logic intensive as it is required to implement several image processing functions on the received image. Within the AMD Versal adaptive SoC architecture, the programmable logic is ideal for implementation of the FEE functionality. The Network On Chip (NOC) enables the image processing chain to be able to store and retrieve images from DDR4 memory, which is often required as part of the image processing chain. Configuration of the FEE functionality, image sensor and image processing chain can be performed by the processor subsystem R5F processors which operate in lockstep.

DHU functionality requires the image to be accessed, compressed and stored within a non-volatile mass memory. Hand over of images from the FEE function to the DHU function can take place using the DDR4 memory and accessed over the NoC. The first stage of the DHU process is lossless compression, and this can be implemented within the programmable logic. Once the image is compressed, the processor subsystem A72 processors can be used to manage the compressed image within the file system in the attached non-volatile memory. The A72 processors can also perform the communication with the satellite computer, retrieving and transmitting the selected images to the satellite computer upon request. The DHU, when not performing image retrieval and storage during active imaging operation, can leverage the processor subsytem and AI engines to retrieve images from the non-volatile memory and use AI algorithms running in the AI engines to classify the image based on elements such as cloud / weather coverage, etc. This enables the DHU to be able to provide an indication of image quality on its stored images to the ground station to ensure the best quality images are downloaded as a priority.

The final element of functionality that can be implemented is the TCU. A typical TCU will provide several analog interface channels that are used to monitor the temperature. Typically, this will use a Platinum Resistance Thermometer (PRT) that requires an analog processing chain that culminates in an analogue to digital convertor. As this is a low frequency control loop, the conversion rate of the ADC is not high speed. As such, commonly used interfaces such as I2C or SPI are used for the ADC. Similarly, the heater drive from the TCU uses a pulse width modulation associated with an analog drive circuit to provide the required delivery power to heat the telescope.

Within the AMD Versal device, the R5F processors can be leveraged to run the PID control loops as required by the update / control cycle (typically 5 to 15 seconds). Interfaces to the external analog input and output domains can leverage the integrated IO peripherals and the programmable logic as necessary for interfacing.

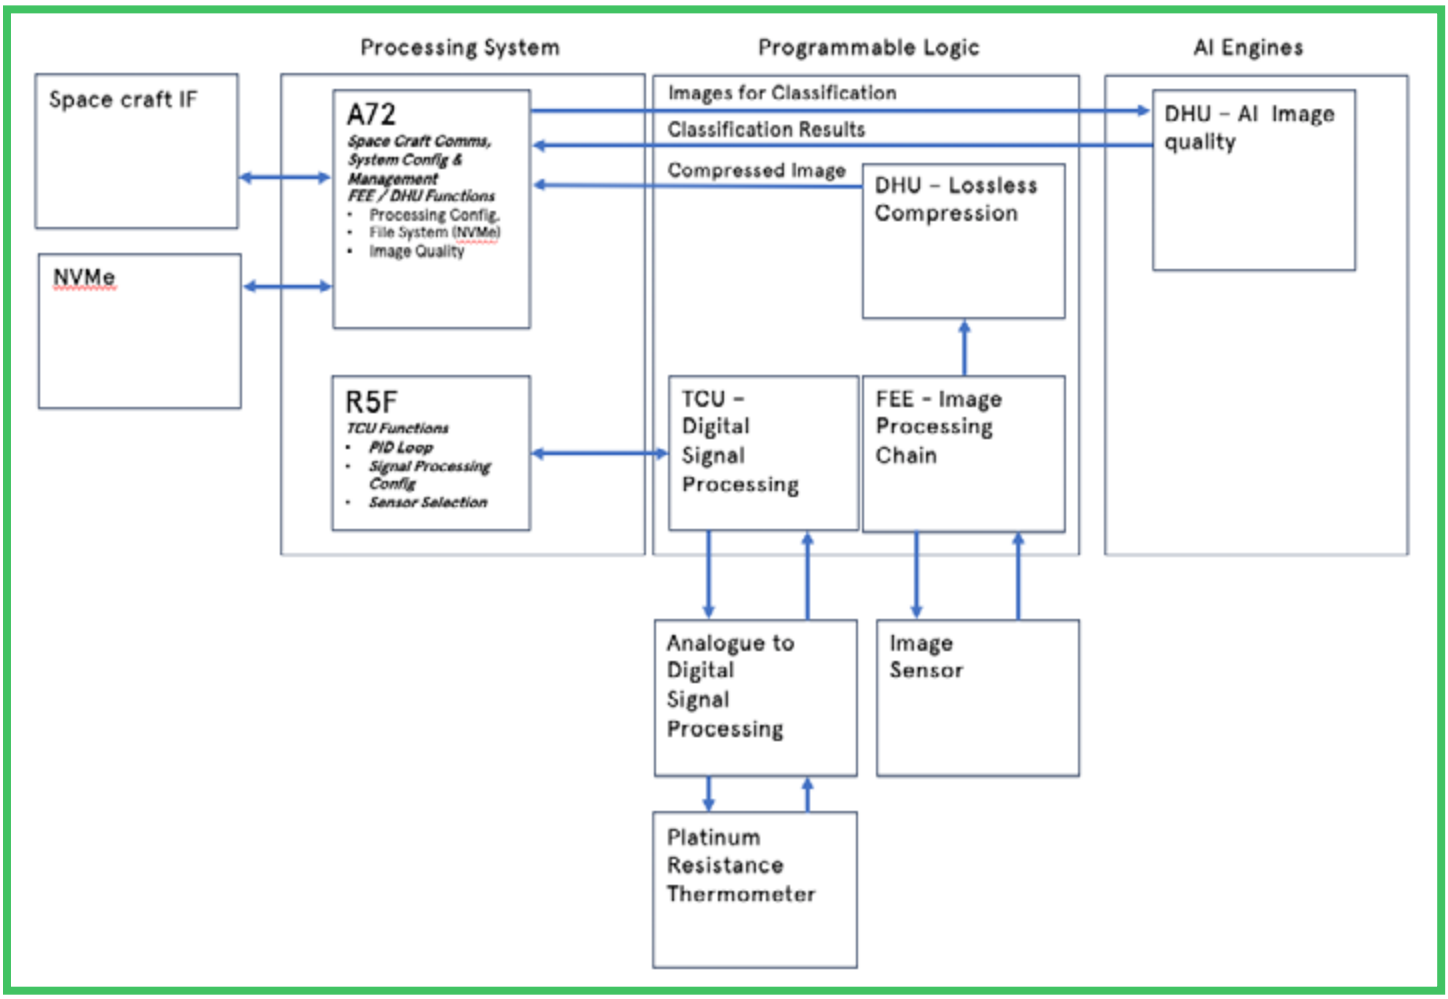

The AMD XQR Versal adaptive SoC therefore enables developers to integrate several different satellite systems into a single system with the processing performed within a single chip solution. At the satellite level, this provides developers with significant benefits to the mass, power, and reliability budgets while also offering potential cost savings.

CONCLUSION

Space plays an increasingly important part in our everyday lives. Developers of space-based systems are faced with increasing performance challenges while still achieving tight power, interfacing and reliability requirements. FPGAs and adaptive SoCs provide developers with a platform and capability to be able to achieve these demanding requirements. FPGAs such as the AMD Kintex UltraScale enable developers to leverage the highly parallel programmable logic to provide a high performance, deterministic solutions. Thanks to the flexibility of the FPGA IO, any to any interfacing can be achieved. Sequential processing within the FPGA can be implemented by leveraging soft core Triple Modular Redundant MicroBlaze solutions within the FPGA fabric.

When it comes to higher performance and system integration, the AMD Versal adaptive SoCs enable developers to achieve high performance parallel processing. The hard processing capabilities and AI engines also enable developers to be able to implement significant system integration. Such integration not only increases performance, but also provides significant system benefits, including reduction of power, and increased reliability.

AMD, the AMD Arrow logo, MicroBlaze, Kintex, Kria, UltraScale, Versal, Vivado, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used herein are for identification purposes and may be trademarks of their respective owners.

Copyright © 2024 Avnet, Inc. The AV logo are registered trademarks of Avnet, Inc.
All other brands are the property of their respective owners. A24_1331_AK

Avnet
2211 S. 47th Street
Phoenix, AZ 85034
1-800-332-8638
avnet.com