Solar Orbiter Will Process Data Onboard Using Xilinx FPGAs

by Tobias Lange
Design Engineer
Braunschweig University of Technology (Germany), Institute of Computer and Network Engineering (IDA)
tobias.lange@tu-bs.de

Holger Michel
Design Engineer
Braunschweig University of Technology, IDA

Björn Fiethe
Senior Engineer
Braunschweig University of Technology, IDA

Harald Michalik
Professor
Braunschweig University of Technology, IDA

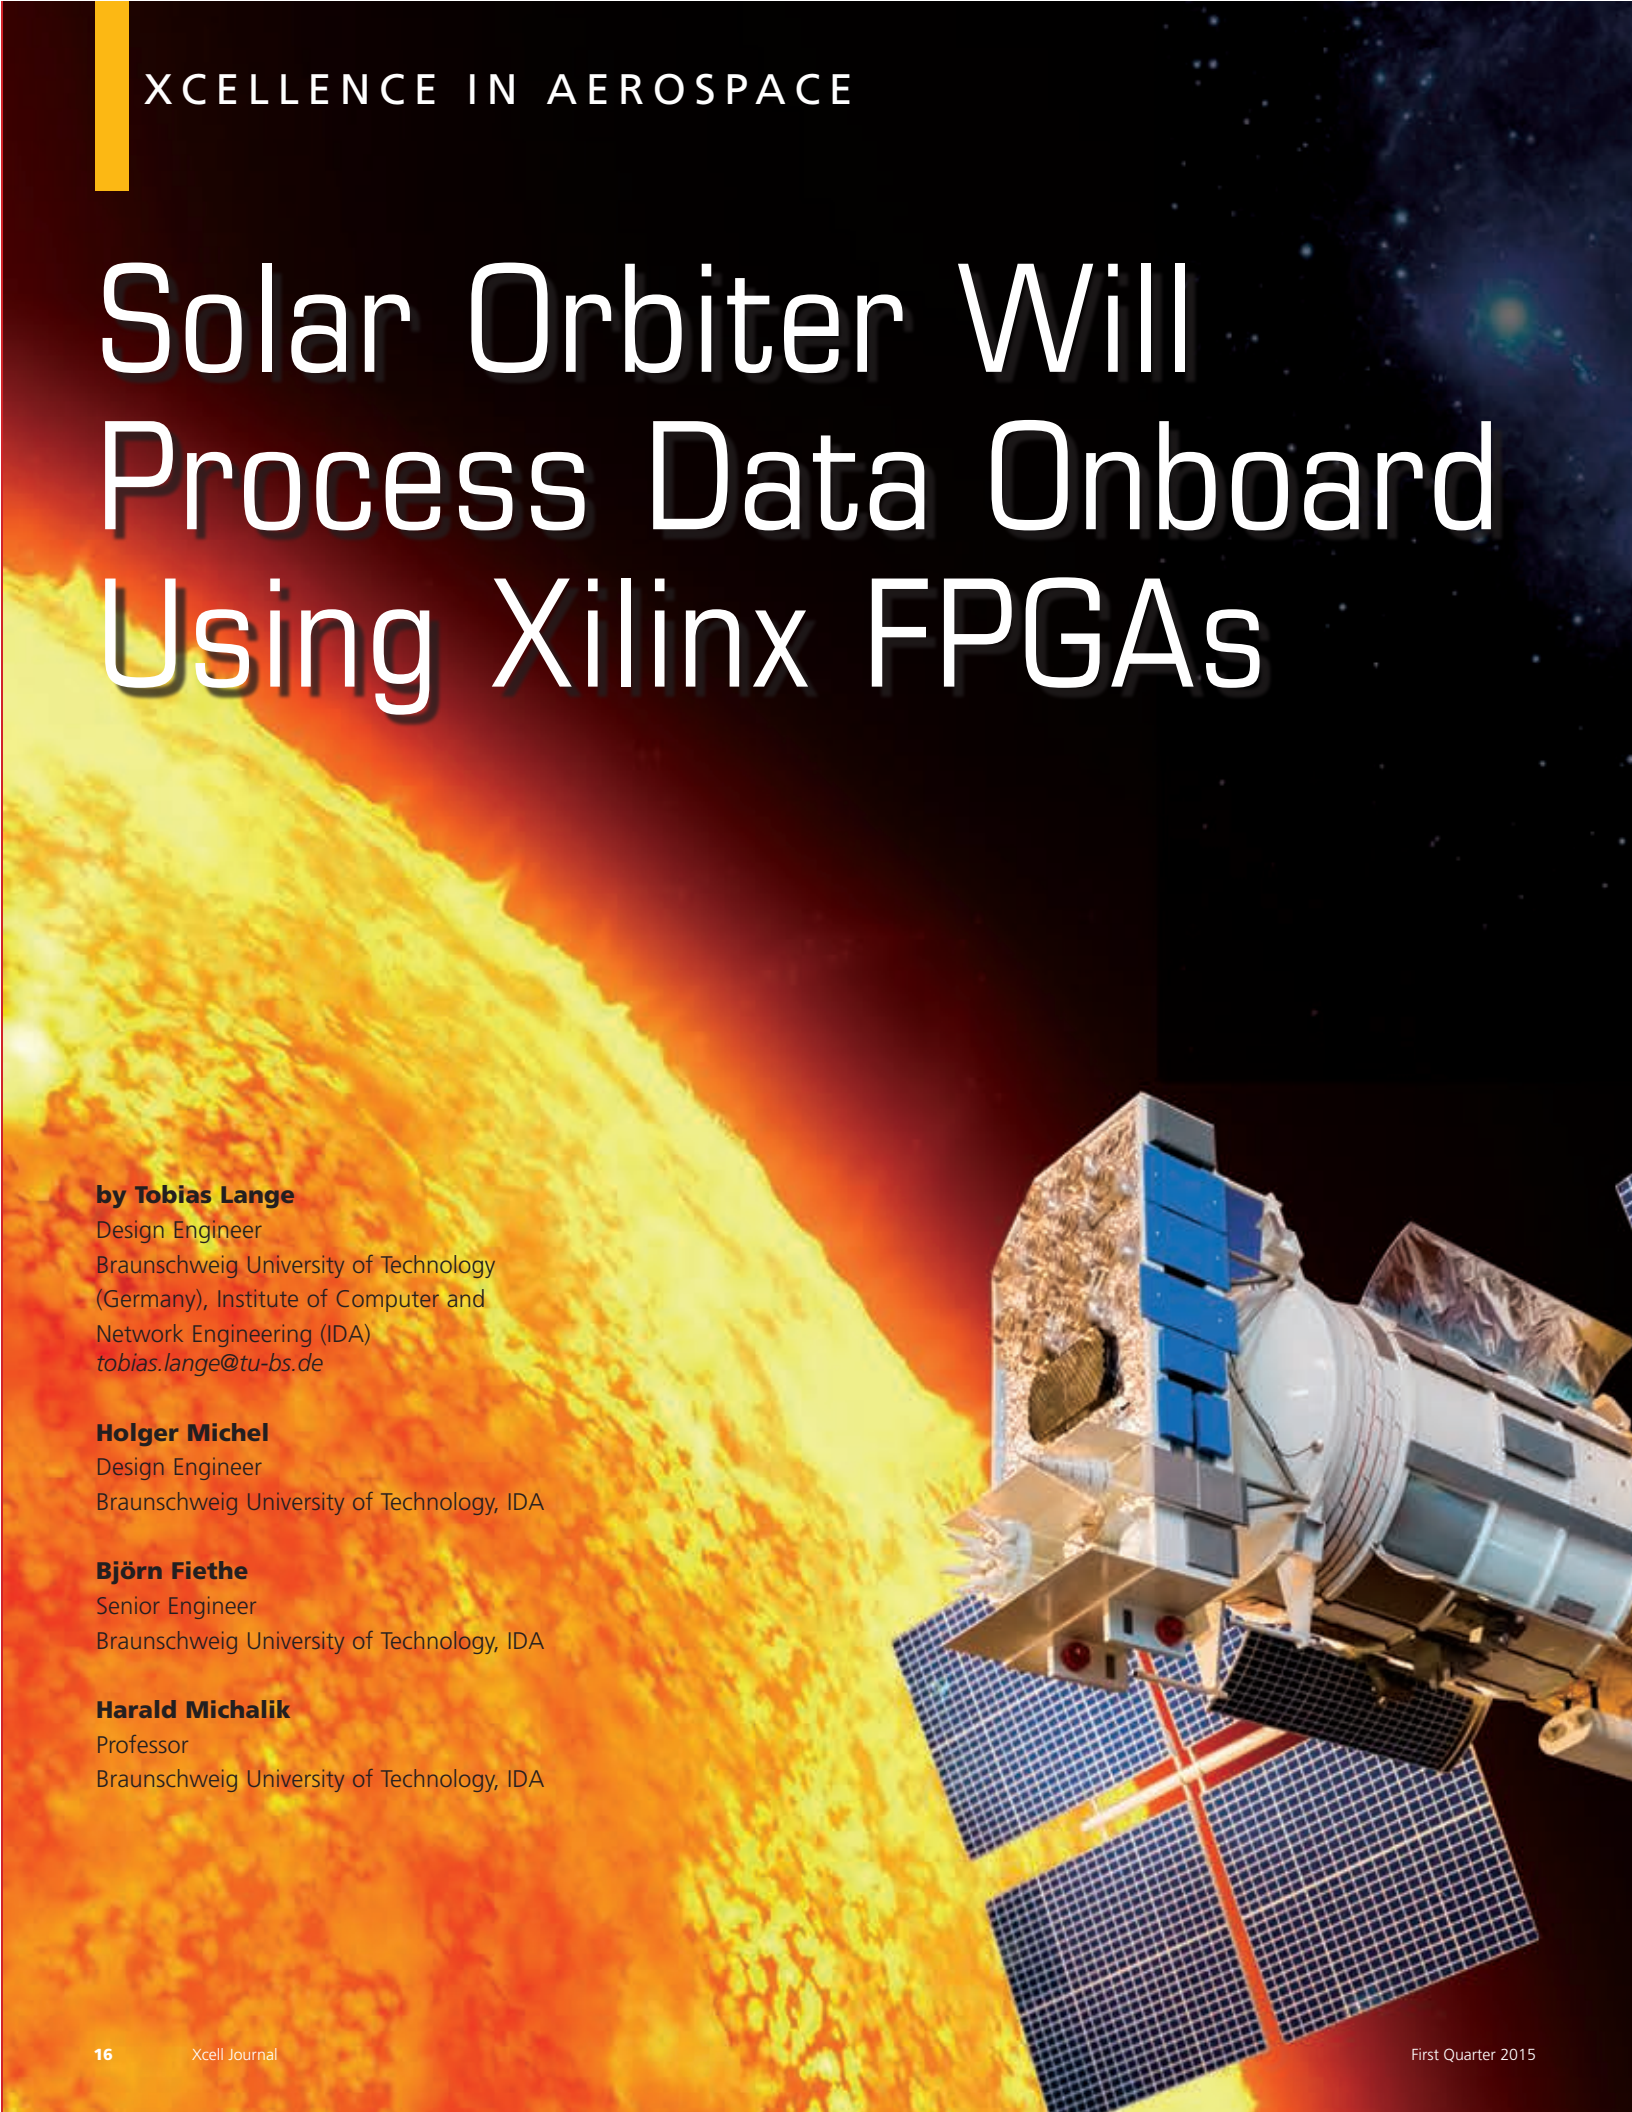

State-of-the-art remote sensing instruments on spacecrafts deliver vast amounts of high-resolution image data. For classical Earth observation missions, scientists typically evaluate the collected data after reception on the ground. While deep-space missions also have to cope with high imaging data rates, the telemetry rate, on the other hand, is very limited.

One demanding example is the Polarimetric and Helioseismic Imager (PHI) instrument, which has been selected as part of the scientific payload for the European Space Agency's Solar Orbiter mission, due to launch in 2017. The PHI instrument, developed mainly at the Max Planck Institute for Solar System Research (MPS) in Germany, will provide maps of the continuum intensity, magnetic-field vector and line-of-sight velocity in the solar photosphere.

Because of the high amount of captured data and the limited downlink capability, extracting scientific parameters onboard the spacecraft will reduce the data volume dramatically. As a result, scientists will be able to take a closer look into the solar photosphere.

To cope with these onboard processing demands, Xilinx® SRAM-based FPGAs with high gate counts offer an attractive solution. Our team at the Braunschweig University of Technology in Germany already has a history with Xilinx FPGAs in active space missions. We have used these devices for classical image-data compression in the processing units of the Venus Express Monitoring Camera (VMC) and the Dawn Framing Camera (DawnFC), both now in successful operation for several years. For the Solar Orbiter PHI processing unit, we decided to use two space-grade Virtex®-4 FPGAs, which will be reconfigured during flight.

Before going into the details of how the FPGAs will streamline data capture on this important mission, let's take a closer look at the PHI itself and examine how it operates.

THE SOLAR ORBITER PHI

The PHI instrument acquires sets of images from an active-pixel sensor. A fil-

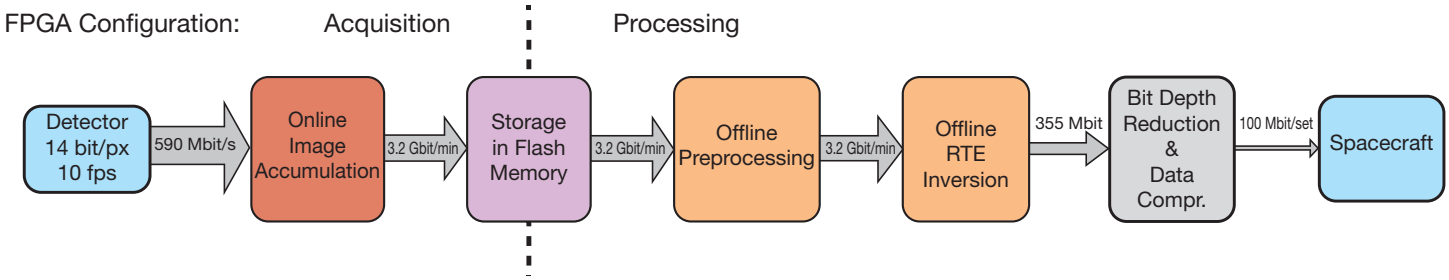

Figure 1 – Polarimetric and Helioseismic Imager (PHI) data-processing pipeline

ter wheel in the optical path applies different wavelength and polarization settings to these images. By preprocessing the captured image data (for example, dark- and flat-field correction) and performing a compute-intensive inversion of the radiative transform equation (RTE), it's possible to calculate magnetic-field vectors from pixel data. Together with standard data compression, this approach will reduce the amount of data from 3.2 Gbits to 100 Mbits per data set. This is a factor of 64 compared with the raw-data input

The processing flow of the PHI can be divided into two modes of operation (Figure 1). Changing between these two modes is perfectly applicable for in-flight reconfiguration of the Virtex FPGAs. Here's how the process works.

During the acquisition phase (the reddish box at left in Figure 1), the detector provides images with a resolution of 2,048 x 2,048 pixels. The acquisition FPGA will accumulate a set of multiple images at different filter settings and directly store them in a large array of NAND flash memory. To reduce residual spacecraft jitter, a dedicated controller for an image-stabilization system will run simultaneously.

After the data acquisition, we reconfigure the two Virtex-4 FPGAs with a preprocessing core and an RTE core (the orange boxes in Figure 1). The preprocessing core retrieves the previously stored data from the flash memory and performs dark- and flat-field correction, addition, multiplication and convolution of frames. Subsequently, the RTE core computes the inversion of the radiative transfer equation.

The FPGA design of the RTE inversion

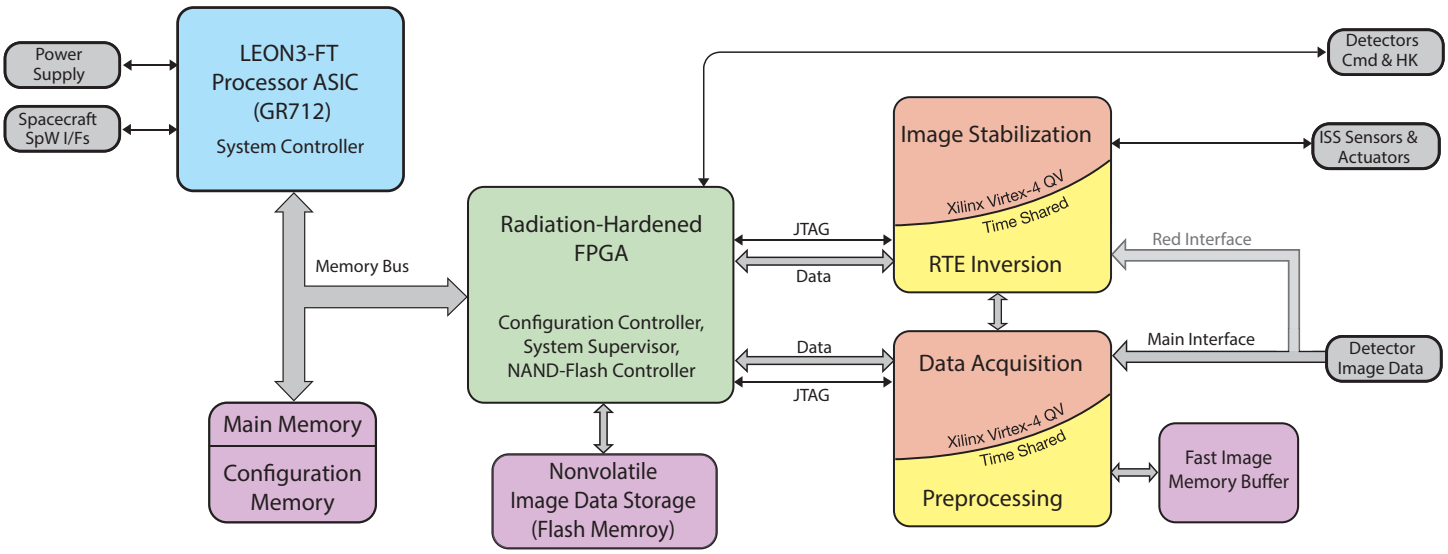

Figure 2 – Architecture of the PHI data-processing unit

Figure 3 – Qualification board with soldered daisychain device

is a contribution of the Instituto de Astrofísica de Andalucía in Granada, Spain. The design for the controller of the image-stabilization system was developed at the University of Barcelona.

ARCHITECTURE OF THE PHI PROCESSING UNIT

Figure 2 shows the architecture of the data-processing unit. For communication to the spacecraft and the system controller, we use a dedicated and highly reliable GR712 LEON3-FT processor ASIC running at a clock frequency of 50 MHz. The CPU has its own 2-Gbit SDRAM memory and is also connected to 1 Gbit of nonvolatile NOR flash memory that stores software and FPGA bitstream configuration files. For the image acquisition, image stabilization, preprocessing and RTE inversion, we use two Virtex-4QV FPGAs. The possibility of in-flight reconfiguration allows us to effectively utilize these two devices in a time-shared manner. This scheme reduces mass, volume and the power consumption of the platform, which are very important factors for deep-space missions.

A one-time-programmable radiation-hardened FPGA connects the LEON3 system controller and the two Virtex-4 devices. Furthermore, this same FPGA functions as a system supervisor. It provides I/O signals and interfaces to control parts of the hardware and the external power supply, and to communicate with sensors and actuators. Two JTAG interfaces allow this FPGA to write and read back the configuration bitstreams of the two Virtex-4 devices.

To store the large amount of image data, we designed a memory board based on an array of NAND flash devices with a total capacity of 4 Tbits. To address this set of memories, located on a separate board, we developed a NAND flash controller that is also placed in the system-supervisor FPGA. To cope with the relatively slow data rate between the NAND flash array and the processing FPGAs, the data acquisition and preprocessing rely on a fast, external buffer memory. A dedicated network connects the system-controller FPGA with the two Virtex-4 FPGAs off-chip and the NAND-flash memory controller with the processing cores on-chip.

DEALING WITH RADIATION EFFECTS

The Xilinx Virtex-4QV is a radiation-tolerant FPGA, which means that the device will not suffer physical damage through radiation effects. Nevertheless, bit upsets can occur and the design has to mitigate them. Radiation can affect SRAM-based

FPGAs in two layers: the configuration layer and the application layer.

Bit upsets in the configuration layer will mainly alter routing logic and combinational functions. One way to repair errors introduced into this layer is to overwrite the configuration SRAM in certain time intervals, a technique known as scrubbing. We optimized the scrubbing process by doing a read back on the bitstream and we reconfigure a certain configuration frame only when an upset is detected.

Radiation effects induced into the application layer will result in faults of the control logic, for example stuck state machines, or simply wrong values in the data path. We will mitigate the upsets on this layer by using triple-modular redundancy (TMR) and error detection and correction (EDAC) techniques.

For a successful mitigation of upsets in the FPGA design, it's crucial to create mechanisms to protect both

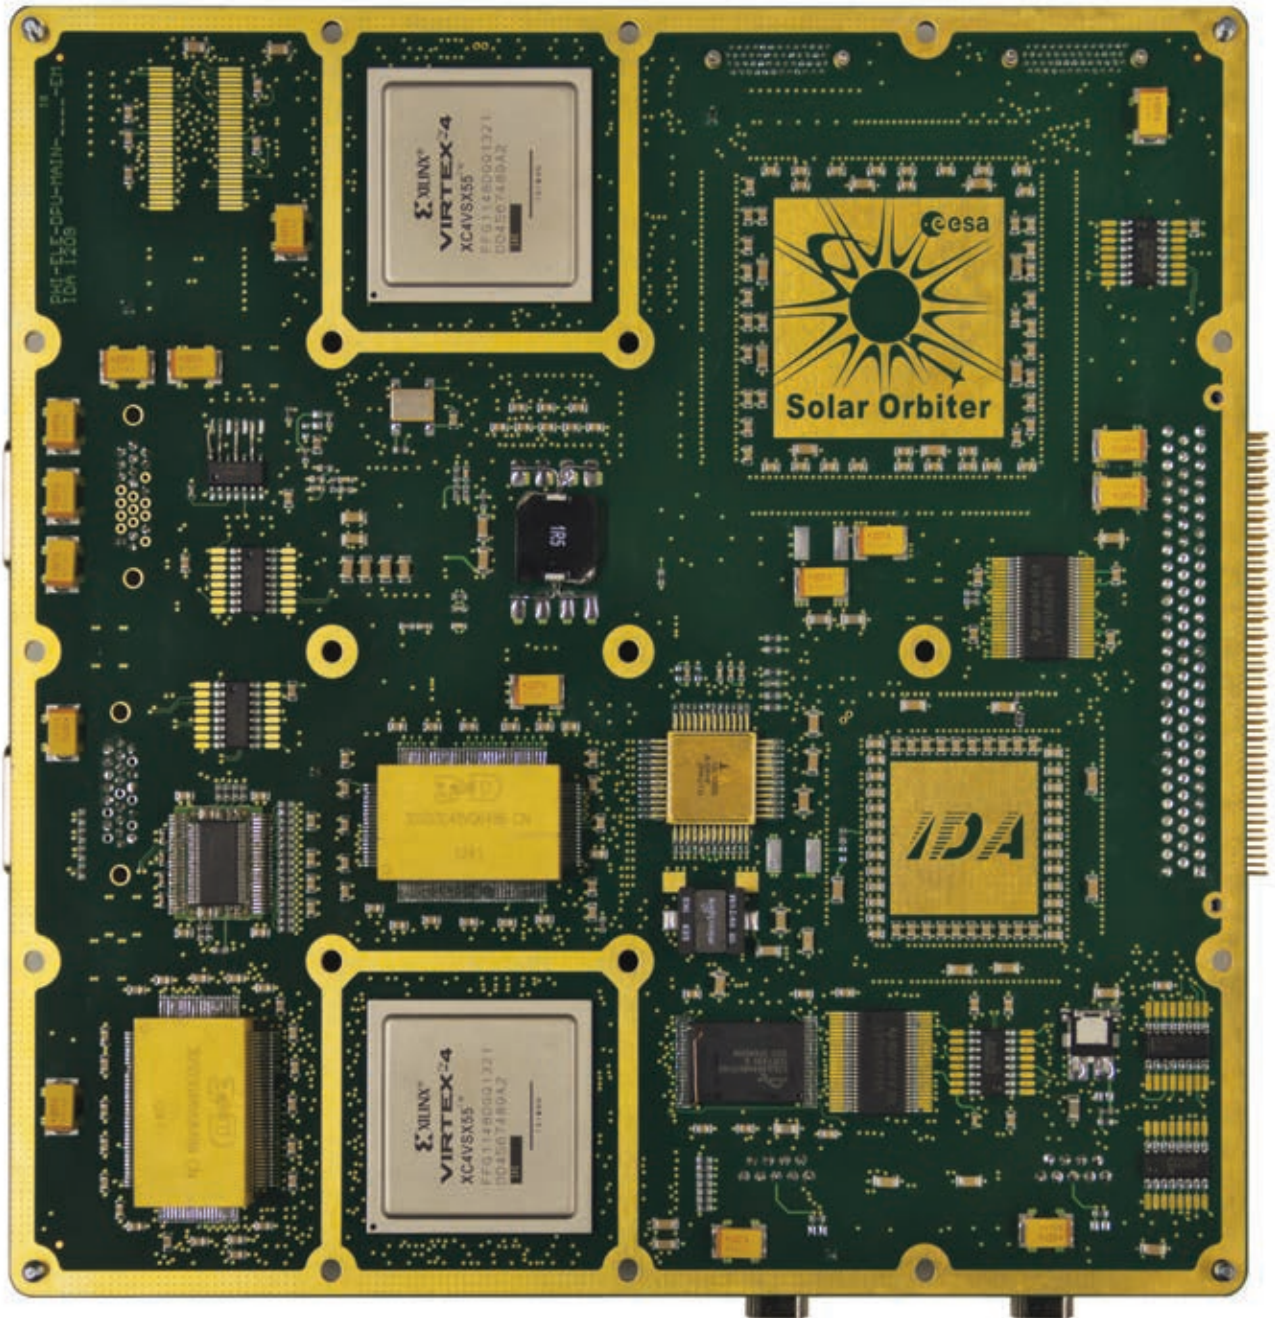

Figure 4 – Bottom side of the PHI data-processing unit engineering model

A high gate count and the ability for in-flight reconfiguration of the Virtex-4 devices made it possible for us to develop a compact and high-performance data-processing platform with reduced size, mass and power consumption.

the configuration layer and the application layer. Integrating only one mitigation scheme alone, scrubbing or TMR, would not be sufficient.

ASSEMBLY QUALIFICATION FOR THE VIRTEX-4QV

Unlike commercial versions of Virtex-4 FPGAs, which come in a flip-chip BGA package, the space-grade Virtex-4QV devices are delivered in a ceramic package. Maintaining a compatible footprint, these devices are equipped with solder columns. When we decided to use Virtex-4QV parts in 2012, no qualified process manufacturer was available in Europe to assemble these CF1140 packages. For this reason we had to start a mission-specific package assembly qualification.

For this purpose, we assembled three representative qualification boards with overall six CF1140 daisy-chain devices (Figure 3). After dedicated shock and vibration tests, we started a thermal-cycling test with close monitoring of resistances of the daisy-chain packages. Before and after each test step, optical and X-ray inspection of the devices proved that no critical physical damage had occurred. We are just finishing the qualification process by means of a destructive micro-sectioning of one PCB.

CURRENT STATUS AND OUTLOOK

After defining a basic architecture for our design including error mitigation and the qualification of the Virtex-4 assembly, we started to work on a prototype of the data-processing unit based on commercial parts. This engineering model fits the final 20 x 20-cm shape of the electronic housing and is already in operation without major problems. It has a mass of 550 grams (without the NAND-flash memory board) and consumes less than 15 watts. The bottom side of this model, equipped with the two Virtex FPGAs, is shown in Figure 4. Currently, we are focusing on finishing the qualification model of our board, equipped with qualified parts.

In summary, the high gate count and the ability for in-flight reconfiguration of the Virtex-4 devices made it possible for us to develop a compact and high-performance data-processing platform with reduced size, mass and power consumption. The data flow of the acquisition and processing perfectly suits a system with two reconfigurable FPGAs.

This system is a first step in bringing in-flight reconfiguration technology to deep-space missions. Electronics for the space industry are usually a couple of years behind commercial technology. Today, the Xilinx Zynq®-7000 All Programmable SoC family integrates SRAM-based FPGA technology with multiple processor cores into a single system-on-chip. In coming years it will be of interest for us to see if these types of SoC solutions will also adapt to the space industry's needs.